

# HP StorageWorks

## Fabric OS 5.0.0 procedures

### user guide

Part number: AA-RW1PA-TE  
First edition: May 2005



## **Legal and notice information**

© Copyright 2005 Hewlett-Packard Development Company, L.P.

© Copyright 2005, Brocade Communications Systems, Incorporated.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information is provided "as is" without warranty of any kind and is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Windows is a U.S. registered trademark of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Linux® is a U.S. registered trademark of Linus Torvalds.

Java™ is a U.S. trademark of Sun Microsystems, Inc.

Fabric OS 5.0.0 procedures user guide

# Contents

<b>About this guide . . . . .</b>	<b>9</b>
Intended audience . . . . .	9
Related documentation . . . . .	9
Document conventions and symbols . . . . .	10
HP technical support . . . . .	11
HP Storage web site . . . . .	11
HP authorized reseller . . . . .	11
<b>1 Introducing Fabric OS CLI procedures . . . . .</b>	<b>13</b>
Changes to this guide for OS v5.0.0 . . . . .	13
Brocade 4Gb SAN Switch for HP p-Class BladeSystem . . . . .	15
About procedural differences . . . . .	15
Scope and references . . . . .	16
About the CLI . . . . .	16
Help information . . . . .	17
Displaying command help . . . . .	17
Displaying additional help topics . . . . .	17
<b>2 Performing basic configuration tasks . . . . .</b>	<b>19</b>
Connecting to the CLI . . . . .	19
Setting the IP address . . . . .	21
Setting the default account passwords . . . . .	21
Setting the date and time . . . . .	22
Maintaining licensed features . . . . .	25
Customizing the switch name . . . . .	27
Customizing the chassis name . . . . .	28
Disabling and enabling a switch . . . . .	28
Disabling and enabling a port . . . . .	29
Activating Ports on Demand . . . . .	30
Making basic connections . . . . .	30
Connecting to devices . . . . .	31
Connecting to other switches . . . . .	31
Working with domain IDs . . . . .	31
Linking through a gateway . . . . .	32
Checking status . . . . .	33
Tracking and controlling switch changes . . . . .	35
<b>3 Configuring standard security features . . . . .</b>	<b>39</b>
Ensuring network security . . . . .	39
Configuring the telnet interface . . . . .	40
Blocking listeners . . . . .	41
Accessing switches and fabrics . . . . .	42
Creating and maintaining user-defined accounts . . . . .	43
To display account information . . . . .	43
To create a user-defined account . . . . .	43
To delete a user-defined account . . . . .	44
To change account parameters . . . . .	44
To recover user-defined accounts . . . . .	45
Changing an account password . . . . .	45
Setting up RADIUS AAA service . . . . .	46

Configuring the RADIUS server . . . . .	48
Windows 2000 . . . . .	48
Linux . . . . .	50
Configuring the switch . . . . .	51
Enabling and disabling local authentication . . . . .	53
Configuring for the SSL protocol . . . . .	54
Browser and Java™ support . . . . .	54
Summary of SSL procedures . . . . .	54
Choosing a CA . . . . .	55
Generating a public/private key . . . . .	55
Generating and storing a CSR . . . . .	56
Obtaining certificates . . . . .	57
Installing a switch certificate . . . . .	57
Activating a switch certificate . . . . .	58
Configuring the browser . . . . .	58
Installing a root certificate to the Java Plug-in . . . . .	59
Displaying and deleting certificates . . . . .	59
Troubleshooting certificates . . . . .	60
Configuring for SNMP . . . . .	60
Setting the security level . . . . .	61
Using the snmpconfig command . . . . .	61
Using legacy commands for SNMPv1 . . . . .	64
Configuring secure file copy . . . . .	69
Setting the boot PROM password . . . . .	70
With a recovery string . . . . .	70
Without a recovery string . . . . .	72
Recovering forgotten passwords . . . . .	73

## **4 Maintaining configurations and firmware . . . . .** 75

Maintaining configurations . . . . .	75
Displaying configuration settings . . . . .	75
Backing up a configuration . . . . .	75
Restoring a configuration . . . . .	76
Downloading configurations across a fabric . . . . .	77
Editing configuration files . . . . .	78
Printing hard copies of switch information . . . . .	78
Maintaining firmware . . . . .	78
Obtaining and unzipping firmware . . . . .	78
Checking connected switches . . . . .	79
About the download process . . . . .	79
Effects of firmware changes on accounts and passwords . . . . .	80
Considerations for downgrading firmware . . . . .	80
Upgrading HP StorageWorks SAN switches . . . . .	80
Summary of the upgrade process . . . . .	81
SAN Switch upgrade procedure . . . . .	81
Upgrading the Core Switch 2/64 and the SAN Director 2/128 . . . . .	83
Summary of the upgrade process . . . . .	83
Core Switch 2/64 and SAN Director 2/128 upgrade procedure . . . . .	84
Troubleshooting firmware downloads . . . . .	86

## **5 Configuring the Core Switch 2/64 and the SAN Director 2/128. . . . .** 89

Identifying ports . . . . .	89
By slot and port number . . . . .	90
By port area ID . . . . .	90
Basic card management . . . . .	90
Powering port cards on and off . . . . .	90
Disabling and enabling cards . . . . .	91
Conserving power . . . . .	91
Setting chassis configurations . . . . .	92

Obtaining slot information . . . . .	92
Configuring a new SAN Director 2/128 with two domains . . . . .	93
Converting an installed SAN Director 2/128 to support two domains . . . . .	94
Combining Core Switch 2/64 and SAN Director 2/128 cards in one chassis . . . . .	95
Setting the card beacon mode . . . . .	98
<b>6 Routing traffic . . . . .</b>	<b>99</b>
About routing policies . . . . .	99
Specifying the routing policy . . . . .	99
Assigning a static route . . . . .	100
Specifying frame order delivery . . . . .	100
Using dynamic load sharing . . . . .	101
Viewing routing path information . . . . .	102
Viewing routing information along a path . . . . .	104
<b>7 Administering extended fabrics . . . . .</b>	<b>107</b>
About extended link buffer allocation . . . . .	107
SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128 . . . . .	107
SAN Switch 4/32 . . . . .	107
Fabric considerations . . . . .	108
Choosing an extended ISL mode . . . . .	108
Configuring an extended ISL . . . . .	109
Trunking over distance . . . . .	111
<b>8 Administering ISL trunking . . . . .</b>	<b>113</b>
Standard trunking criteria . . . . .	113
Fabric considerations . . . . .	114
Initializing trunking on ports . . . . .	115
Monitoring traffic . . . . .	115
Enabling and disabling ISL trunking . . . . .	116
Setting port speeds . . . . .	117
Displaying trunking information . . . . .	119
Trunking over extended fabrics . . . . .	119
Troubleshooting trunking problems . . . . .	120
Listing link characteristics . . . . .	120
Recognizing buffer underallocation . . . . .	120
<b>9 Administering advanced zoning . . . . .</b>	<b>123</b>
Zoning terminology . . . . .	123
Zoning concepts . . . . .	124
Zone types . . . . .	124
Zone objects . . . . .	126
Zone aliases . . . . .	126
Zone configurations . . . . .	126
Zoning enforcement . . . . .	127
Software-enforced zoning . . . . .	127
Hardware-enforced zoning . . . . .	127
Rules for configuring zones . . . . .	130
Creating and managing zone aliases . . . . .	131
Creating and maintaining zones . . . . .	133
Creating and modifying zoning configurations . . . . .	134
Managing zoning configurations in a fabric . . . . .	137
Adding a new switch or fabric . . . . .	137
Splitting a fabric . . . . .	139
Using zoning to administer security . . . . .	139
Resolving zone conflicts . . . . .	140

<b>10 Administering advanced performance monitoring . . . . .</b>	<b>143</b>
Displaying and clearing the CRC error count . . . . .	145
Monitoring end-to-end performance . . . . .	145
Adding end-to-end monitors . . . . .	146
Setting a mask for end-to-end monitors . . . . .	148
Deleting end-to-end monitors . . . . .	149
Monitoring filter-based performance . . . . .	149
Adding standard filter-based monitors . . . . .	150
Adding custom filter-based monitors . . . . .	150
Deleting filter-based monitors . . . . .	152
Monitoring ISL performance . . . . .	152
Monitoring trunks . . . . .	152
Displaying monitor counters . . . . .	153
Clearing monitor counters . . . . .	155
Saving and restoring monitor configurations . . . . .	156
Collecting performance data . . . . .	156
<b>11 Configuring the distributed management server . . . . .</b>	<b>157</b>
Enabling and disabling the platform services . . . . .	157
Controlling access . . . . .	158
Configuring the server database . . . . .	161
Controlling topology discovery . . . . .	162
<b>12 Working with diagnostic features . . . . .</b>	<b>165</b>
Viewing power-on self test . . . . .	165
Viewing switch status . . . . .	166
Viewing port information . . . . .	168
Viewing equipment status . . . . .	171
Viewing the system message log . . . . .	172
Viewing the port log . . . . .	173
Configuring for syslogd . . . . .	175
Configuring the Host . . . . .	176
Configuring the switch . . . . .	176
Viewing and saving diagnostic information . . . . .	177
Setting up automatic trace dump transfers . . . . .	178
<b>13 Troubleshooting . . . . .</b>	<b>181</b>
Most common problem areas . . . . .	182
Gathering information for technical support . . . . .	182
Analyzing connection problems . . . . .	184
Restoring a segmented fabric . . . . .	186
Correcting zoning setup issues . . . . .	187
Recognizing MQ-WRITE errors . . . . .	189
Correcting I2C bus errors . . . . .	190
Correcting device login issues . . . . .	191
Identifying media-related issues . . . . .	194
Correcting link failures . . . . .	196
Correcting marginal links . . . . .	198
Inaccurate information in the system message log . . . . .	200
Recognizing the port initialization and FCP auto discovery process . . . . .	200
<b>A Configuring the PID format . . . . .</b>	<b>203</b>
About PIDs and PID binding . . . . .	203
Summary of PID formats . . . . .	204
Impact of changing the fabric PID format . . . . .	204
Host reboots . . . . .	205
Static PID mapping errors . . . . .	205
Changes to configuration data . . . . .	205

Selecting a PID format . . . . .	206
Evaluating the fabric . . . . .	207
Planning the update procedure . . . . .	209
Online update . . . . .	209
Offline update . . . . .	210
Hybrid update . . . . .	210
Changing to core PID format . . . . .	211
Changing to extended edge PID format . . . . .	212
Performing PID format changes . . . . .	213
Basic procedure . . . . .	214
HP-UX procedure . . . . .	215
AIX procedure . . . . .	216
Swapping port area IDs . . . . .	217
<b>B Configuring interoperability mode . . . . .</b>	<b>219</b>
Vendor switch requirements . . . . .	219
HP StorageWorks switch requirements . . . . .	219
Supported HP StorageWorks features . . . . .	220
Unsupported HP StorageWorks features . . . . .	220
Configuration recommendations . . . . .	221
Configuration restrictions . . . . .	221
Zoning restrictions . . . . .	221
Zone name restrictions . . . . .	222
Enabling and disabling interoperability mode . . . . .	222
<b>C Using Remote Switch . . . . .</b>	<b>225</b>
<b>D Understanding legacy password behavior . . . . .</b>	<b>229</b>
Password management information . . . . .	229
Password prompting behaviors . . . . .	232
Password migration during firmware changes . . . . .	233
Password recovery options . . . . .	235
<b>E Zone merging scenarios . . . . .</b>	<b>237</b>
<b>F Upgrading firmware in single mode . . . . .</b>	<b>239</b>
<b>Glossary . . . . .</b>	<b>243</b>
<b>Index . . . . .</b>	<b>267</b>

## Figures

1 Hardware-enforced non-overlapping zones . . . . .	129
2 Hardware-enforced overlapping zones . . . . .	129
3 Zoning with hardware assist (mixed port and WWN zones) . . . . .	130
4 Overlapping hardware-enforced zoning with soft porting . . . . .	130
5 Setting end-to-end monitors on a port . . . . .	147
6 Proper placement of end-to-end performance monitors . . . . .	147
7 Mask positions for end-to-end monitors . . . . .	148
8 Configure command on a switch running Fabric OS 3.1.2 . . . . .	213
9 Configure command on a switch running Fabric OS 4.2.0 and later . . . . .	213

## Tables

1 Document conventions . . . . .	10
2 Brocade 4Gb SAN Switch for HP p-Class BladeSystem . . . . .	15
3 Conversion from UTC to local time . . . . .	24

4	Standard ISL modes . . . . .	31
5	Blocked listener applications . . . . .	42
6	Access details . . . . .	42
7	SSL certificate files . . . . .	55
8	Recommended CAs . . . . .	55
9	Commands for displaying and deleting SSL certificates . . . . .	59
10	SSL messages and actions . . . . .	60
11	Recommended firmware . . . . .	79
12	Effect of firmware on accounts and passwords . . . . .	80
13	Supported options . . . . .	92
14	Header fields . . . . .	92
15	Extended ISL modes . . . . .	108
16	SAN Switch 2/8V, 2/16V, 2/32, Core Switch 2/64, and SAN Director 2/128 . . . . .	109
17	SAN Switch 4/32 . . . . .	109
18	Types of zoning . . . . .	124
19	Approaches to fabric-based zoning . . . . .	125
20	Enforcing hardware zoning . . . . .	128
21	Zoning database limitations . . . . .	137
22	Considerations for zoning architecture . . . . .	140
23	Advanced performance monitoring commands . . . . .	144
24	Commands to add filter-based monitors . . . . .	150
25	Predefined values at offset 0 . . . . .	151
26	Error summary description . . . . .	170
27	Commands for port log management . . . . .	174
28	Fabric OS to UNIX message severities . . . . .	176
29	Common troubleshooting problems and tools . . . . .	182
30	Types of zone discrepancies . . . . .	187
31	Commands for debugging zoning . . . . .	188
32	Component test descriptions . . . . .	194
33	Switch component tests . . . . .	195
34	SwitchShow output and suggested action . . . . .	198
35	Effects of PID format changes on configurations . . . . .	205
36	PID format recommendations for adding new switches . . . . .	207
37	PID format and management interface names . . . . .	211
38	Earliest Fabric OS versions for extended edge PID format . . . . .	212
39	Account and password characteristics matrix . . . . .	229
40	Password prompting matrix . . . . .	232
41	Password migration behavior during firmware upgrade and downgrade . . . . .	233
42	Password recovery options . . . . .	235
43	Zone merging scenarios . . . . .	237

# About this guide

This document provides information to assist fabric administrators in using the web-based graphical user interface to monitor and modify their HP StorageWorks switch fabrics.

This preface discusses the following topics:

- [Intended audience](#), page 9
- [Related documentation](#), page 9
- [Document conventions and symbols](#), page 10
- [HP technical support](#), page 11

## Intended audience

This book is intended for use by those responsible for monitoring and modifying their HP StorageWorks switch fabric.

## Related documentation

Documentation, including white papers and best practices documents, is available via the HP website. Please go to:

<http://www.hp.com/country/us/eng/prodserv/storage.html>

To access 4.x related documents:

1. Locate the **Networked storage** section of the web page.
2. Under **Networked storage**, go to the **By type** subsection.
3. Click **SAN infrastructure**. The SAN infrastructure page displays.
4. Locate the **Fibre Channel Switches** section.

Locate the **B-Series Fabric** subsection, and then go to the appropriate subsection, such as **Enterprise Class** for the SAN Director 2/128.

To access 4.x documents (such as this document), select the appropriate product, for example **SAN Director 2/128 & 2/128 Power Pack** or **Core Switch 2/64 & Core Switch 2/64 Power Pack**.

The switch overview page displays.

5. Go to the **Product information** section, located on the far right side of the web page.
6. Click **Technical documents**.
7. Follow the onscreen instructions to download the applicable documents.

# Document conventions and symbols

**Table 1** Document conventions

Convention	Element
Medium blue text: <a href="#">Figure 1</a>	Cross-reference links and e-mail addresses
Medium blue, underlined text ( <a href="http://www.hp.com">http://www.hp.com</a> )	Web site addresses
<b>Bold font</b>	<ul style="list-style-type: none"><li>Key names</li><li>Text typed into a GUI element, such as into a box</li><li>GUI elements that are clicked or selected, such as menu and list items, buttons, and check boxes</li></ul>
<i>Italics</i> font	Text emphasis
Monospace font	<ul style="list-style-type: none"><li>File and directory names</li><li>System output</li><li>Code</li><li>Text typed at the command-line</li></ul>
<i>Monospace italic font</i>	<ul style="list-style-type: none"><li>Code variables</li><li>Command-line variables</li></ul>
<b>Monospace, bold font</b>	Emphasis of file and directory names, system output, code, and text typed at the command-line



**WARNING!** Indicates that failure to follow directions could result in bodily harm or death.



**CAUTION:** Indicates that failure to follow directions could result in damage to equipment or data.



**IMPORTANT:** Provides clarifying information or specific instructions.



**NOTE:** Provides additional information.



**TIP:** Provides helpful hints and shortcuts.

# HP technical support

Telephone numbers for worldwide technical support are listed on the following HP web site:  
<http://www.hp.com/support/>. From this web site, select the country of origin.



**NOTE:** For continuous quality improvement, calls may be recorded or monitored.

Obtain the following information before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

## HP Storage web site

The HP web site has the latest information on this product, as well as the latest drivers. Access storage at:  
<http://www.hp.com/country/us/eng/prodserv/storage.html>. From this web site, select the appropriate product or solution.

## HP authorized reseller

For the name of your nearest HP authorized reseller:

- In the United States, call 1-800-345-1518.
- Elsewhere, visit <http://www.hp.com> and click **Contact HP** to find locations and telephone numbers.



# 1 Introducing Fabric OS CLI procedures

This guide contains procedures for configuring and managing an HP StorageWorks Storage Area Network (SAN) using the Fabric OS Command Line Interface (CLI). This chapter consists of the following sections:

- [Changes to this guide for OS v5.0.0](#), page 13
- [About procedural differences](#), page 15
- [Scope and references](#), page 16
- [About the CLI](#), page 16
- [Help information](#), page 17

The guide applies to the following HP products:

- HP StorageWorks Switches: 1-GB switches, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32.

These switches contain a fixed number of ports (they are called *fixed-port switches*). The SAN Switch 4/32 allows you to license and activate extra fixed ports with the Ports on Demand feature.

- Core Switch 2/64 and SAN Director 2/128.

These switches can contain a variable number of ports, which you install by plugging port cards into the director chassis.

## Changes to this guide for OS v5.0.0

The following changes are new to v5.0.0 and are not included elsewhere in this guide.

- Add “Brocade 4Gb SAN Switch for HP p-Class BladeSystem” everywhere that the HP StorageWorks SAN Switch 4/32 is mentioned, except as specified in the following sections.
  - On page 43, in the section “Creating and maintaining a user-defined account,” change the following definition item:

**-r rolename** Specifies the role: either admin or user in nonsecure mode; admin, user, or nonfcsadmin in secure mode.  
to:  
**-r rolename** Specifies the role: either admin or user in nonsecure mode or admin, user, switchAdmin, or nonfcsadmin in secure mode.
  - On page 44, in the section “To change account parameters,” change the following definition item:

**-r rolename** Specifies the role: either admin or user in nonsecure mode; admin, user, or nonfcsadmin in secure mode.  
to:  
**-r rolename** Specifies the role: either admin or user in nonsecure mode; admin, user, switchAdmin, or nonfcsadmin in secure mode.

- On page 107, in the section “SAN Switch 4/32” add the following:

For the Brocade 4Gb SAN Switch for HP p-Class BladeSystem, each port group contains four ports and buffer credits are shared among all ports on the switch.

- On page 108, in the section “Choosing an extended ISL mode,” add the following:

## Brocade 4Gb SAN Switch for HP p-Class BladeSystem

The number of ports that can be configured at various distances is summarized in [Table 2](#).

**Table 2** Brocade 4Gb SAN Switch for HP p-Class BladeSystem

Speed (Gbit/sec)	Number of ports allowed at distance (km)			
	1	2	3	4
1	286	154	110	88
2	143	77	55	44
4	71.5	38.5	27.5	22

- On page 146, in the section “Adding end-to-end monitors,” change the following paragraph:

The HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, SAN Director 2/128, allow up to eight end-to-end monitors allow up to eight end-to-end monitors.

to:

The HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, SAN Director 2/128, and Brocade 4Gb SAN Switch for HP p-Class BladeSystem allow up to eight end-to-end monitors.

- On page 149, in the section “Monitoring filter-based performance,” change the following paragraph:

For HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128, the maximum number of filters is eight per port, in any combination of standard filters and user-defined filters.

to:

For HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, SAN Director 2/128, and Brocade 4Gb SAN Switch for HP p-Class BladeSystem, the maximum number of filters is eight per port, in any combination of standard filters and user-defined filters.

- On page 219, in the section “HP StorageWorks switch requirements,” add the following:

Brocade 4Gb SAN Switch for HP p-Class BladeSystem running Fabric OS 5.0.0 or later.

## About procedural differences

As a result of the differences between fixed-port and variable-port devices, procedures sometimes differ between HP StorageWorks switch models. Also, because the domain architecture of the Core Switch 2/64 differs from that of the SAN Director 2/128, there are sometimes procedural differences between these two.

When procedures or parts of procedures apply to some models but not others, this guide identifies the specifics for each model. For example, a number of procedures that apply only to variable-port devices are found in [“Configuring the Core Switch 2/64 and the SAN Director 2/128”](#) on page 89. Procedures that apply only to the SAN Switch 4/32 are labeled as such.



**NOTE:** When command examples in this guide show user input enclosed in quotation marks, the quotation marks are required for versions earlier than v4.0.0. They are optional in later versions, unless specifically called for in the procedures.

## Scope and references

Although many different software and hardware configurations are tested and supported by HP, documenting all possible configurations and scenarios is beyond the scope of this document. In some cases, earlier releases are highlighted to present considerations for interoperating with them.

The hardware reference manuals for HP StorageWorks products describe how to power up devices and set their IP addresses. After the IP address is set, you can use the CLI procedures contained in this guide.

This guide provides only the level of detail required to perform the procedures. If you need more information about the commands used in the procedures, refer to online help or to the *HP StorageWorks Fabric OS 4.x command reference guide*.

There are several access methods that you can use to configure a switch. These are listed with their respective documents:

- For Advanced Web Tools procedures, refer to the *HP StorageWorks Fabric OS 4.x Advanced Web Tools user guide*.
- For Fabric Manager procedures, refer to the *HP StorageWorks Fabric OS 4.4.x Fabric Manager user guide*.
- For third-party application procedures, refer to the third-party API documentation.

## About the CLI

Fabric OS CLI is the complete fabric management tool for HP SANs that enables you to:

- Access the full range of Fabric OS features, based on license keys.
- Configure, monitor, dynamically provision, and manage every aspect of the SAN.
- Configure and manage the HP fabric on multiple efficient levels.
- Identify, isolate, and manage SAN events across every switch in the fabric.
- Manage switch licenses.
- Perform fabric stamping.

To manage a switch using telnet, SNMP, and HP Advanced Web Tools, the switch must be connected to a network through the switch Ethernet port (out of band) or from the Fibre Channel (in band). The switch must be configured with an IP address to allow for the network connection. Refer to the hardware manual for your switch for information on physically connecting to the switch.

You can access switches from different connections, such as Advanced Web Tools, CLI, and API. When these connections are simultaneous, changes from one connection may not be updated to the other, and some modifications may be lost. When simultaneous connections are used, make sure that you do not overwrite the work of another connection.

In a mixed fabric containing switches running various Fabric OS versions, you should use the latest-model switches running the most recent release for the primary management tasks. The principal management

access should be set to the core switches in the fabric. For example, to run Secure Fabric OS, use the latest-model switch as the primary FCS, the location to perform zoning tasks, and the time server.

A number of management tasks are designed to make fabric-level changes; for example, zoning commands make changes that affect the entire fabric. When executing fabric-level configuration tasks, allow time for the changes to propagate across the fabric before executing any subsequent tasks. For a large fabric, it may be take a few minutes.

## Help information

Each Fabric OS command provides Help information that explains the command function, its possible operands, its level in the command hierarchy, and additional pertinent information.

### Displaying command help

To display help information:

1. Connect to the switch and log in as admin.
2. To display a list of all command help topics, issue the `help` command with no arguments.
3. To display help for a specific command, enter `help command`, where `command` is the name of the command for which you need information, as shown in the following example:

```
switch:admin> help configure
Administrative Commands                                     configure(1m)
NAME
    configure - change system configuration settings
SYNOPSIS
    configure
AVAILABILITY
    admin
DESCRIPTION
    This command changes some system configuration settings,
    including:
        o Arbitrated loop settings
        o Switch fabric settings
        o System services settings
        o Virtual channel settings
    (output truncated)
```

### Displaying additional help topics

The following commands provide help files for specific topics:

<code>diagHelp</code>	Diagnostic help information
<code>fwHelp</code>	Fabric Watch help information
<code>licenseHelp</code>	License help information
<code>perfHelp</code>	Performance Monitoring help information
<code>routeHelp</code>	Routing help information
<code>trackChangesHelp</code>	Track Changes help information
<code>zoneHelp</code>	Zoning help information



## 2 Performing basic configuration tasks

This chapter contains procedures for performing basic switch configuration tasks using the Fabric OS CLI and contains the following sections:

- [Connecting to the CLI](#), page 19
- [Setting the IP address](#), page 21
- [Setting the default account passwords](#), page 21
- [Setting the date and time](#), page 22
- [Maintaining licensed features](#), page 25
- [Customizing the switch name](#), page 27
- [Customizing the chassis name](#), page 28
- [Disabling and enabling a switch](#), page 28
- [Disabling and enabling a port](#), page 29
- [Activating Ports on Demand](#), page 30
- [Making basic connections](#), page 30
- [Working with domain IDs](#), page 31
- [Linking through a gateway](#), page 32
- [Checking status](#), page 33
- [Tracking and controlling switch changes](#), page 35

### Connecting to the CLI

You can connect to the CLI either through a telnet connection or through the serial port. To connect with telnet:

1. Verify that the switch is connected to the IP network through the RJ-45 Ethernet port.  
Switches in the fabric that are not connected via Ethernet can be managed through switches that are using IP over Fibre Channel. The embedded port must have an assigned IP address.
2. Open a telnet connection to the switch.  
The login prompt is displayed when the telnet connection finds the switch in the network.  
For the Core Switch 2/64 and SAN Director 2/128, enter the logical switch name (sw0 or sw1).
3. Enter the account ID (defaults are user or admin) at the login prompt.
4. Enter the password. The default password is password.  
If you have not changed the system passwords from the default, you are prompted to change them.
5. Enter the new system passwords, or press **Ctrl+c** to skip the password prompts.

6. Verify that the login was successful. The prompt displays the switch name and user ID to which you are connected:

```
login: admin  
password: xxxxxxxx  
switch:admin>
```

Consider the following for telnet connections:

- Never change the IP address of the switch while two telnet sessions are active; if you do, your next attempt to log in fails. To recover, gain access to the switch by one of these methods:
  - Use Advanced Web Tools and perform a fast boot. When the switch comes up, the telnet quota is cleared. For instructions on performing a fast boot with Advanced Web Tools, refer to the *HP StorageWorks Fabric OS 4.x Advanced Web Tools user guide*.
  - If you have the required privileges, you can connect through the serial port, log in as root, and use operating system commands to identify and kill the telnet processes without disrupting the fabric.
- For admin level accounts, Fabric OS limits the number of simultaneous telnet sessions per switch to two. For more details on session limits, see "[Configuring the telnet interface](#)" on page 40 and "[Creating and maintaining user-defined accounts](#)" on page 43.

To connect through the serial port:

- Connect the serial cable to the serial port on the switch and to an RS-232 serial port on the workstation.

If the serial port on the workstation is RJ-45 instead of RS-232, remove the adapter on the end of the serial cable and insert the exposed RJ-45 connector into the RJ-45 serial port on the workstation.

- Open a terminal emulator application (such as HyperTerminal on a PC, or TERM, TIP, or Kermit in a UNIX® environment), and configure the application as follows:
  - In a Windows® environment:

Parameter	Value
Bits per second	9600
Databits	8
Parity	None
Stop bits	1
Flow control	None

- In a UNIX environment, enter the following string at the prompt: tip /dev/ttys -9600  
If ttys is already in use, you can use ttysa (enter tip /dev/ttysa -9600).

Consider the following for serial connections:

- Some procedures require that you connect through the serial port; for example, setting the IP address or setting the boot PROM password.
- If secure mode is enabled, connect through the serial port of the primary FCS switch.
- For the Core Switch 2/64 and SAN Director 2/128, you can connect to CP0 or CP1 using either of the two serial ports.

# Setting the IP address

You must connect through the serial port to set the IP address (see "To connect through the serial port:" on page 20). After connecting, use the `ipaddrset` command to set the IP address.



**CAUTION:** The use of IP address 0.0.0.0 is not supported. Do not use this address.

Fabric OS v2.6.0, v3.1.0, and v4.0.0 supports Classless Inter-Domain Routing (CIDR).

# Setting the default account passwords

For each logical switch (domain), there are admin and user default access accounts. These accounts designate different levels of authorization—called *roles*—for using the system:

- The *admin* level is for administrative use.
- The *user* level is for nonadministrative use, such as monitoring system activity.

Two accounts—`factory` and `root`—are reserved for development and manufacturing. You can change their passwords, which is optional, but you should not use these accounts under normal circumstances.

For the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, SAN Switch 4/32, and SAN Director 2/128 (default configuration with one domain), there is one set of default access accounts.

For the Core Switch 2/64 and SAN Director 2/128, configured with two domains, each logical switch has its own set of default access accounts. The default account names and passwords are the same for both of the logical switches.

You can also create up to 15 additional accounts per logical switch and designate their roles as either admin or user. See the procedures in "Creating and maintaining user-defined accounts" on page 43.

For large enterprises, Fabric OS supports RADIUS services, as described in "Setting up RADIUS AAA service" on page 46.

In addition to the account access passwords, each switch can set a boot PROM password. For greater security, HP recommends that you set this password to protect system boot parameters from unauthorized access. See "Setting the boot PROM password" on page 70.

Each of the default access accounts has an associated password. The first time you connect to a Fabric OS switch you are prompted to change these default account passwords.

If you do not change the default passwords, you are prompted to do so at each subsequent login until all system passwords have been changed from the default values. Thereafter, use the `passwd` command to change passwords.

For more background information on passwords, see "Changing an account password" on page 45.

To change the default passwords at login:

1. Connect to the switch and log in as admin.
2. At each of the `Enter new password` prompts, either enter a new password or skip the prompt. Skip a prompt by pressing `Enter`. You can bypass all further prompts by pressing `Ctrl+c`.

Although the `root` and `factory` accounts are not meant for general use, you should change their passwords if prompted to do so and save the passwords in case they are needed for recovery purposes.

Passwords can be from 8 to 40 characters long and must begin with an alphabetic character. They can include numerals, the dot (.), and the underscore (\_). They are case sensitive, and they are not displayed when you enter them on the command line. You cannot reuse the default passwords.



**NOTE:** Record the passwords exactly as entered and store them in a secure place; recovering passwords requires significant effort and fabric downtime.

Example:

```
login: admin
Password:
Please change your passwords now.
Use Control-C to exit or press 'Enter' key to proceed.
for user - root
Changing password for root
Enter new password: *****
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
Please change your passwords now.
for user - factory
Changing password for factory
Enter new password: *****
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
Please change your passwords now.
for user - admin
Changing password for admin
Enter new password: *****
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
Please change your passwords now.
for user - user
Changing password for user
Enter new password: *****
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
switch:admin>
```

## Setting the date and time

Switches maintain the current date and time in nonvolatile memory. Date and time are used for logging events. Switch operation does not depend on the date and time; a switch with an incorrect date and time value still functions properly. However, because the date and time are used for logging, you should set them correctly.



**NOTE:** The date and tsclockserver commands are disabled when the security feature is enabled. With security enabled you can view the current date setting only on the primary FCS switch.

To set the date and time:

1. Connect to the switch and log in as admin.
2. Enter the `date` command at the command line using the following syntax:

```
date "MMDhmmYY"
```

The values represent the following:

- *MM* is the month; valid values are 01 through 12.
- *DD* is the date; valid values are 01 through 31.
- *hh* is the hour; valid values are 00 through 23.
- *mm* is minutes; valid values are 00 through 59.
- *YY* is the year; valid values are 00 through 99 (values greater than 69 are interpreted as 1970 through 1999, and values less than 70 are interpreted as 2000 through 2069).



**NOTE:** The date function does not support daylight savings time or time zones, so such changes must be reset manually.

Example:

```
switch:admin> date
Fri May  5 21:50:00 UTC 1989
switch:admin>
switch:admin> date "0624165203"
Tue Jun 24 16:52:30 UTC 2003
switch:admin>
```

You can synchronize the local time of the principal or primary Fabric Configuration Server (FCS) switch to an external NTP server.

To synchronize local time with an external source:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
tsclockserver ipaddr
```

where *ipaddr* is the IP address of the NTP server, which the switch must be able to access. This operand is optional; by default this value is LOCL, which uses the local clock of the principal or primary switch as the clock server.

Example:

```
switch:admin> tsclockserver
LOCL
switch:admin> tsclockserver "132.163.135.131"
switch:admin> tsclockserver
132.163.135.131
switch:admin>
```

HP recommends that you synchronize time with an external NTP server, as described on [page 23](#). If you cannot do so, use the next procedure.

To set the time zone:

1. Connect to the switch and log in as admin.
2. Issue the `tstimezone` command as follows:

```
tstimezone [houroffset [, minuteoffset]]  
• For Pacific Standard Time enter tsTimeZone -8, 0  
• For Central Standard Time enter tsTimeZone -6, 0  
• For Eastern Standard Time enter tsTimeZone -5, 0
```

The default time zone for switches is Universal Time Conversion (UTC), which is 8 hours ahead of (later than) Pacific Standard Time (PST). For additional time zone conversions, see [Table 3](#).

The parameters do not apply if the time zone of the switch has already been changed from the default (8 hours ahead of PST).

Refer to the `tstimezone` command in the *HP StorageWorks Fabric OS 4.x command reference guide* for more detailed information about the command parameters.

3. Repeat the procedure on all switches for which the Time Zone needs to be set. This needs to be done only once, because the value is written to nonvolatile memory.

For U.S. time zones, use [Table 3](#) to determine the correct parameter for the `tstimezone` command.

**Table 3** Conversion from UTC to local time

Local time	<code>tstimezone</code> conversion parameter
Atlantic Standard	-4,0
Atlantic Daylight	-3,0
Eastern Standard	-5,0
Eastern Daylight	-4,0
Central Standard	-6,0
Central Daylight	-5,0
Mountain Standard	-7,0
Mountain Daylight	-6,0
Pacific Standard	-8,0
Pacific Daylight	-7,0
Alaskan Standard	-9,0
Alaskan Daylight	-8,0
Hawaiian Standard	-10,0

# Maintaining licensed features

Feature licenses might be part the licensed Paper Pack supplied with switch software, or you can purchase licenses separately from your switch vendor, who will provide you with keys to unlock the features. License keys are provided on a per-chassis basis, so for products that support multiple logical switches (domains), a license key applies to all domains within the chassis.

To unlock a licensed feature, you can either use the license key in the Paper Pack supplied with switch software, or launch an internet browser and go to the HP licensing web site at <http://webkey.external.hp.com/welcome.asp>. Click **Generate a license key** and follow the instructions to generate the key.



**NOTE:** You need the following items for each chassis to be licensed:

- Transaction key, which is in the Paper Pack supplied with switch software. Or, when you purchase a license, your switch vendor gives you a transaction key to be used for obtaining a license key.
- License ID; to see a switch License ID, use the `licenseIDShow` command.

To unlock a licensed feature:

1. If you already have a license key, go to [step 10](#).

If you do not have a key, launch an Internet browser and go to the HP web site:  
<http://www.hp.com/country/us/eng/prodserv/storage.html>.

2. Click **products**.
3. Click **Software Products**.
4. In the **Related Links** panel on the right side of the page, select **Software License Keys**.

The Software License Keys instruction page appears.

5. If you want to generate a single license key, select **Generate 1 license key**.  
If you want to generate multiple license keys, select **Batch Generation of Licenses**.  
The Software License Key instruction page appears.
6. Enter the requested information in the required fields.

When generating multiple license keys, enter the worldwide names and transaction keys in the table at the bottom of the screen. If you need additional rows in the table, select **Add More Rows**.

7. Click **Next**.  
A verification screen appears.
8. Verify that the information appears correctly.  
Click **Submit** if the information displayed is correct. If the information is incorrect, Click **Previous** and change the information.
9. After the information is corrected, click **Submit**.  
An information screen displays the license keys.  
You also receive an e-mail with the keys and installation instructions.

**10.** Activate and verify the license as follows:

- a. Connect to the switch and log in as admin.
- b. Activate the license using the **licenseadd** command. For example:

```
switch:admin> licenseadd "key"
```

The license key is case sensitive and must be entered exactly as given. The quotation marks are optional.

For the Core Switch 2/64 and SAN Director 2/128, the licenses are effective on both CPs and on all logical switches.

- c. Verify that the license was added by issuing the **licenseshow** command. The licensed features currently installed on the switch are listed. If the feature is not listed, issue the **licenseadd** command again.
- d. Some features may require additional configuration , or you might need to disable and reenable the switch to make them operational; see the feature documentation for details.

**Example**

```
switch:admin> licenseshow
SbeSdQdQySyriTeJ:
Web license
Zoning license
Fabric license
Remote Switch license
Extended Fabric license
Fabric Watch license
Performance Monitor license
Trunking license
Security license
SbbebdQS9QTscfcB:
Ports on Demand license - additional 8 port upgrade
SbbebdQS9QTcgfcz:
Ports on Demand license - additional 8 port upgrade
```

To remove a licensed feature:

1. Connect to the switch and log in as admin.
2. Issue the **licenseshow** command to display the active licenses.
3. Remove the license key using the **licenseremove** command. For example:

```
switch:admin> licenseremove "key"
```

The license key is case sensitive and must be entered exactly as given. The quotation marks are optional.

After removing a license key, the optionally licensed feature is disabled when the switch is rebooted or when a switch disable or enable is performed. For the Core Switch 2/64 and SAN Director 2/128, reboot both the primary and the secondary CP cards to ensure that HA features remain synchronized.

4. Issue the **licenseshow** command to verify that the license is disabled.

Example:

```
switch:admin> licenseshow
bQebzbRdScRfc0iK:
    Web license
    Zoning license
SybbzQQ9edTzcc0X:
    Fabric license
switch:admin> licenseremove "bQebzbRdScRfc0iK"
removing license key "bQebzbRdScRfc0iK"
switch:admin>
```

After a reboot (or switchdisable and switchenable):

Example:

```
switch:admin> licenseshow
SybbzQQ9edTzcc0X:
    Fabric license
switch:admin>
```

If there are no license keys, licenseshow displays No licenses.

## Customizing the switch name

Switches can be identified by IP address, Domain ID, World Wide Name (WWN), or by customized switch names that are unique and meaningful.

Version 4.0.0 and later switch names can be from 1 to 15 characters; they must begin with a letter and may contain letters, numbers, and the underscore character. It is not necessary to use quotation marks.

The default names are:

- swd77 for the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32.
- For the Core Switch 2/64, the two logical switches have different default names. The name swd77 is used for the logical switch containing the port cards in slots 1 through 4; swd76 is used for the logical switch containing the port cards in slots 7 through 10.
- swd77 for the SAN Director 2/128



**NOTE:** Changing the switch name causes a domain address format RSCN to be issued.

To customize the switch name:

1. For the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32, proceed to the next step.

For the Core Switch 2/64 and the SAN Director 2/128, identify the serial console for the active CP. You can do so by issuing the hashow command from any Core Switch 2/64 and SAN Director 2/128 serial console, or by looking for the blue Active LED on the SAN Director 2/128.

2. Connect to the switch and log in as admin.

**3.** For the Core Switch 2/64 and the SAN Director 2/128, proceed to the next step.

For the SAN Director 2/128, if configured for one domain (the default) proceed to the next step. If configured with two domains, proceed as for the Core Switch 2/64.

For the Core Switch 2/64, choose the logical switch that you want to change. Enter the value that corresponds to that logical region:

- Enter 0 to configure logical switch 0 (slot 1 through 4).
- Enter 1 to configure logical switch 1 (slot 7 through 10).

**4.** Issue the **switchname** command at the command line with the following syntax:

**switchname "newname"**

where *newname* is the new name for the switch.

**5.** Record the new switch name for future reference.

**6.** For the Core Switch 2/64 and the SAN Director 2/128 configured with two domains, disconnect from the session and repeat the procedure for the second logical switch.

Example:

```
switch:admin> switchname "switch62"
Committing configuration...
Done.
switch62:admin>
```

---

## Customizing the chassis name

Beginning with Fabric OS v4.4.0, HP recommends that you customize the chassis name for each switch. Some system logs identify switches by chassis names, so if you assign meaningful chassis names in addition to meaningful switch names, logs are more useful.

To change the chassis name:

1. Connect to the switch and log in as admin.
2. Issue the **chassisname** command at the command line with the following syntax:

**chassisname "newname"**

Where *newname* is the new name for the chassis.

Chassis names can be from 1 to 15 characters, must begin with a letter, and can contain letters, numbers, and the underscore character. It is not necessary to use the quotation marks.

3. Record the new chassis name for future reference.

---

## Disabling and enabling a switch

By default, the switch is enabled after power is applied and diagnostics and switch initialization routines have finished. You can disable and reenable it as necessary.

To disable a switch:

1. Connect to the switch and log in as admin.

- Issue the `switchdisable` command at the command line.

All Fibre Channel ports on the switch are taken offline. If the switch was part of a fabric, the fabric reconfigures.

To enable a switch:

- Connect to the switch and log in as admin.
- Issue the `switchenable` command at the command line.

All Fibre Channel ports that passed the POST test are enabled. If the switch has interswitch links to a fabric, it joins the fabric.

---

## Disabling and enabling a port

All licensed ports are enabled by default. You can disable and reenable them as necessary. Ports that you activate with Ports on Demand must be enabled explicitly, as described in "[Activating Ports on Demand](#)" on page 30.

To disable a port:

- Connect to the switch and log in as admin.
- For the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32, issue the following command:

**portdisable** *portnumber*

where *portnumber* is the port number of the port you want to disable.

For the Core Switch 2/64 and the SAN Director 2/128, issue the following command:

**portdisable** *slotnumber/portnumber*

where *slotnumber* and *portnumber* are the slot and port numbers of the port you want to disable.

If the port is connected to another switch, the fabric might reconfigure.

To enable a port:

- Connect to the switch and log in as admin.
- For the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32, issue the following command:

**portenable** *portnumber*

The *portnumber* is the port number of the port you want to enable.

For the Core Switch 2/64 and the SAN Director 2/128, issue the following command:

**portenable** *slotnumber/portnumber*

The *slotnumber* and *portnumber* are the slot and port numbers of the port you want to enable. (Slots are numbered 1 through 4 and 7 through 10, counting from left to right.)

If the port is connected to another switch, the fabric might reconfigure. If the port is connected to one or more devices, these devices become available to the fabric.

If you change port configurations during a switch failover, the ports might become disabled. Reissue the `portenable` command after the failover is complete to bring the ports online.

# Activating Ports on Demand

The Core Switch 4/32 can be purchased with 16 or 32 licensed ports. As your needs increase, you can activate the remaining ports by purchasing and installing the optional HP StorageWorks 8-port upgrade license.

Ports on Demand is ready to be unlocked in the switch firmware. Its license might be part of the licensed Paper Pack supplied with switch software, or you can purchase the license separately from your switch vendor, who will provide you with a key to unlock it.

By default, ports 0 through 15 are enabled on the SAN Switch 4/32. To enable ports 16 through 23, purchase and install an 8-port upgrade license. To enable ports 24 through 31, purchase and install another 8-port upgrade license. The first license key must be already installed before you can use the second license.

You must enable the ports after you have installed the license keys. You can do so without disrupting switch operation using the `portenable` command on each port. Alternatively, you can disable and reenable the switch to activate all ports.

To enable an 8-port upgrade license, you can either use the supplied license key or generate a license key. If you need to generate a key, launch an Internet browser and go to the HP licensing web site at <http://webkey.external.hp.com/welcome.asp>. Click **Generate a license key** and follow the instructions to generate the key.

To enable Ports on Demand:

1. Connect to the switch and log in as admin.
2. Optional: to verify the states of the ports, use the `portshow` command.  
In the `portshow` output, the Licensed field shows whether the port is licensed.
3. Install the HP Ports on Demand licensed product.  
For instructions, see "Maintaining licensed features" on page 25.
4. Use the `portenable` command to enable the ports.
5. Optional: use the `portshow` command to check the newly activated ports.

If you remove an 8-port upgrade license, the licensed ports become disabled after the next platform reboot or the next port deactivation.

# Making basic connections

You can make basic connections to devices and to other switches.

Before connecting a v4.0.0 or later switch to a fabric that contains switches running earlier firmware versions, you must first set the same PID format on all the switches. The presence of different PID formats in a fabric causes fabric segmentation.

For information on PID formats and related procedures, see "Selecting a PID format" on page 206.

For information on configuring the routing of connections, see "Routing traffic" on page 99.

For information on configuring extended interswitch connections, see "Administering extended fabrics" on page 107.

## Connecting to devices

To minimize port logins, power off all devices before connecting them to the switch. For devices that cannot be powered off, first use the `portdisable` command to disable the port on the switch, and then connect the device. When powering the devices back on, wait for each device to complete the fabric login before powering on the next one.

## Connecting to other switches

Refer to the hardware user's guide of your specific switch for interswitch link (ISL) connection and cable management information. [Table 4](#) summarizes the standard ISL modes, which you can configure with the `portcfglongdistance` command. For information on extended ISL modes, which enable longer distance interswitch links, see "[Administering extended fabrics](#)" on page 107.

**Table 4** Standard ISL modes

Mode	Description	Maximum ISL distance (km)	Earliest fabric OS release
LO <sup>1</sup>	Level 0 static mode, the default.	<ul style="list-style-type: none"><li>• 10 km at 1 Gbps</li><li>• 5 km at 2 Gbps</li><li>• 2.5 km at 4 Gbps</li></ul>	All
LE	Level E static mode, supports links beyond 5 km.	10 km at 1, 2, or 4 Gbps	v3.0.0, v4.0.0

1. When you upgrade from Fabric OS v4.0.0 to Fabric OS v4.1.0 or later, all extended ISL ports are set automatically to LO mode.

## Working with domain IDs

Although domain IDs are assigned dynamically when a switch is enabled, you can reset them manually so that you can control the ID number or to resolve a domain ID conflict when you merge fabrics.

If a switch already has a domain ID when it is enabled, and that domain ID conflicts with a switch already in the fabric, the conflict is automatically resolved. The process can take several seconds, during which time traffic is delayed.

The default domain ID for HP StorageWorks switches is 1.

The default domain ID applies to both of the logical switches in Core Switch 2/64 and SAN Director 2/128 switches that are configured for two domains. To prevent domain conflict, you can either disable one of the switches until both are connected to the fabric, then reenable the switches so that unique domain IDs are automatically assigned; or use the procedure "[To set the domain ID:](#)" on page 32 to make the domain IDs unique before connecting the logical switches to the fabric.



**CAUTION:** On switches running Fabric OS v4.0.0 and later, do not use domain ID 0, which is reserved for another purpose. Using this domain ID can cause the switch to reboot continuously. Avoid changing the domain ID on the FCS in secure mode. To minimize down time, change the domain IDs on the other switches in the secure fabric.

To display domain IDs:

1. Connect to a switch and log in as admin.

**2. Issue the fabricshow command.**

Fabric information is displayed, including the domain ID (D\_ID).

Example:

```
switch:admin> fabricshow
Switch ID      Worldwide Name          Enet IP Addr    FC IP Addr    Name
-----
3: fffc43      10:00:00:60:69:10:60:1f  192.168.64.187  0.0.0.0      "sw187"
2: ffc42       10:00:00:60:69:00:05:91  192.168.64.60   192.168.65.60  "sw60"
1:fffc41       10:00:00:60:69:00:02:0b  192.168.64.180  192.168.65.180 > "sw180"
The Fabric has 3 switches
Group ID      Token
-----
0: fffb01     40:05:00:00:10:00:00:60:69:00:00:15
```

The fields in the fabricshow display are:

Switch ID      The switch Domain\_ID and embedded port D\_ID.

Worldwide Name The switch WWN.

Enet IP Addr    The switch Ethernet IP address.

FC IP Addr     The switch FC IP address.

Name            The switch symbolic name. An arrow (>) indicates the principal switch.

To set the domain ID:

**1.** Connect to the switch and log in as admin.

**2.** Issue the switchdisable command to disable the switch.

**3.** Issue the configure command.

**4.** Enter **y** after the Fabric Parameters prompt:

```
Fabric parameters (yes, y, no, n): [no] y
```

**5.** Enter a unique domain ID at the Domain prompt. Use a domain ID value from 1 through 239 for normal operating mode (FCSW compatible). For example:

```
Domain: (1..239) [1] 3
```

**6.** Respond to the remaining prompts (or press **Ctrl+d** to accept the other settings and exit).

**7.** Issue the switchenable command to reenable the switch.

## Linking through a gateway

A gateway merges SANs into a single fabric—by establishing point-to-point E\_Port connectivity between two Fibre Channel switches that are separated by a network—with a protocol such as IP or SONET.

Except for link initialization, gateways are transparent to switches; the gateway simply provides E\_Port connectivity from one switch to another.

By default, switch ports initialize links using the Exchange Link Parameters (ELP) mode 1. However, gateways expect initialization with ELP mode 2. Therefore, to enable two switches to link through a gateway, the ports on both switches must be set for ELP mode 2.

Any number of E\_Ports in a fabric can be configured for gateway links, if these rules are followed:

- All switches in the fabric must be upgraded to Fabric OS v3.1.0 (or later) or v4.1.0 (or later).
- To prevent fabric segmentation, make sure that all switches in the fabric use the core PID format, as described in “[To configure a link through a gateway](#)” on page 33.

- When determining switch count maximums, include the switches connected to both sides of the gateway.
- Extended links (those created using the Extended Fabrics licensed feature) and the security features in Secure Fabric OS are not supported through gateway links.

To configure a link through a gateway:

- If you are not sure that the PID format is consistent across the entire fabric, issue the `configshow` command on all switches to check the PID setting. If necessary, change the PID format on any nonconforming switches as described in “[Configuring the PID format](#)” on page 203.
- Connect to the switch on one end of the gateway and log in as admin.
- Issue the `portcfgislmode` command:

SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32	<code>portcfgislmode port mode</code>
	Specify a <i>port</i> number. Valid values for port number vary, depending on the switch type. The <i>mode</i> operand is required: specify 1 to enable ISL R_RDY mode (gateway link) or specify 0 to disable it.
Core Switch 2/64 and SAN Director 2/128	<code>portcfgislmode slot/port, mode</code>
	Specify a <i>slot/port</i> number pair. Valid values for slot and port number vary depending on the switch type. The <i>mode</i> operand is required: specify 1 to enable ISL R_RDY mode (gateway link) or specify 0 to disable it.

In the following example, slot 2, port 3 is enabled for a gateway link:

```
switch:admin> portcfgislmode 2/3, 1
Committing configuration...done.
ISL R_RDY Mode is enabled for port 3. Please make sure the PID
formats are consistent across the entire fabric.
switch:admin>
```

- Repeat the previous steps for any additional ports to be connected to the gateway.
- Repeat the procedure on the switch at the other end of the gateway.

Refer to the *HP StorageWorks Fabric OS 4.x command reference guide* for more information about the `portcfgislmode` command.

## Checking status

You can check the status of switch operation, HA features, and fabric connectivity:

- Connect to the switch and log in as admin.
- Issue the `switchshow` command at the command line.  
This command displays a switch summary and a port summary.
- Verify that the switch and ports are online.
- Issue the `switchstatusshow` command to further check the status of the switch.

To verify HA features:

HA features provide maximum reliability and nondisruptive replacement of key hardware and software modules. To verify these features, connect to the switch as admin and use any of the following commands:

- `chassisshow` verifies the Field Replaceable Units (FRUs).
- For the Core Switch 2/64 and the SAN Director 2/128:
  - `hashow` verifies that HA is enabled, that the heartbeat is up, and that the HA state is synchronized between the active and standby CP cards.
  - `slotshow` inventories and displays the current status of each slot in the system.

To verify fabric connectivity:

1. Connect to the switch and log in as admin.
2. Issue the `fabricshow` command, which displays a summary of all the switches in the fabric.

Example:

```
switch:admin> fabricshow
Switch ID    Worldwide Name          Enet IP Addr   FC IP Addr   Name
-----+
 1: fffc01 10:00:00:60:69:80:04:5a 192.168.186.61 192.168.68.193 "switch61"
 3: fffc03 10:00:00:60:69:10:9c:29 192.168.186.175 0.0.0.0      "switch175"
 4: fffc04 10:00:00:60:69:12:14:b7 192.168.174.70 0.0.0.0      "switch70"
 5: fffc05 10:00:00:60:69:45:68:04 192.168.144.121 0.0.0.0      "switch121"
 6: fffc06 10:00:00:60:69:00:54:ea 192.168.174.79 192.168.68.197 "switch79"
 7: fffc07 10:00:00:60:69:80:04:5b 192.168.186.62 192.168.68.194 "switch62"
 8: fffc08 10:00:00:60:69:04:11:22 192.168.186.195 0.0.0.0      "switch195"
 9: fffc09 10:00:00:60:69:10:92:04 192.168.189.197 192.168.68.198 "switch197"
10: fffc0a 10:00:00:60:69:50:05:47 192.168.189.181 192.168.68.181 "switch181"
11: fffc0b 10:00:00:60:69:00:54:e9 192.168.174.78 192.168.68.196 "switch78"
15: fffc0f 10:00:00:60:69:30:1e:16 192.168.174.73 0.0.0.0      "switch73"
33: fffc21 10:00:00:60:69:90:02:5e 192.168.144.120 0.0.0.0      "switch120"
44: fffc2c 10:00:00:60:69:c0:06:8d 192.168.144.121 0.0.0.0      "switch121"
97: fffc61 10:00:00:60:69:90:02:ed 192.168.144.123 0.0.0.0      "switch123"
98: fffc62 10:00:00:60:69:90:03:32 192.168.144.122 0.0.0.0      "switch122"

The Fabric has 15 switches

switch:admin>
```

To verify device connectivity:

1. Connect to the switch and log in as admin.
2. Optional: Issue the `switchshow` command to verify that devices, hosts, and storage are connected.
3. Optional: Issue the `nsshow` command to verify that devices, hosts, and storage have successfully registered with the Name Server.
4. Issue the `nsallshow` command, which displays 24-bit Fibre Channel addresses of all devices in the fabric.

Example:

```
switch:admin> nsallshow
{
    010e00 012fef 012fef 030500 030b04 030b08 030b17 030b18
    030b1e 030b1f 040000 050000 050200 050700 050800 050de8
    050def 051700 061c00 071a00 073c00 090d00 0a0200 0a07ca
    0a07cb 0a07cc 0a07cd 0a07ce 0a07d1 0a07d2 0a07d3 0a07d4
    0a07d5 0a07d6 0a07d9 0a07da 0a07dc 0a07e0 0a07e1 0a0f01
    0a0f02 0a0f0f 0a0f10 0a0f1b 0a0f1d 0b2700 0b2e00 0b2fe8
    0b2fef 0f0000 0f0226 0f0233 0f02e4 0f02e8 0f02ef 210e00
    211700 211fe8 211fef 2c0000 2c0300 611000 6114e8 6114ef
    611600 620800 621026 621036 6210e4 6210e8 6210ef 621400
    621500 621700 621a00
75 Nx_Ports in the Fabric }
switch:admin>
```

The number of devices listed should reflect the number of devices that are connected.

## Tracking and controlling switch changes

The Track Changes feature allows you to keep a record of specific changes that may not be considered switch events, but can provide useful information. The output from the Track Changes feature is dumped to the system messages log for the switch. Use the `errdump` or `errshow` command to view the log.

Items in the log created from the Track Changes feature are labeled `TRACK`.

Trackable changes are:

- Successful login
- Unsuccessful login
- Logout
- Configuration file change from task
- Track Changes on
- Track Changes off

An SNMP-TRAP mode can also be enabled; refer to the `trackchangeshelp` command in the *HP StorageWorks Fabric OS 4.x command reference guide*.

For troubleshooting information on the Track Changes feature, see "[Inaccurate information in the system message log](#)" on page 200.

To enable the Track Changes feature:

1. Connect to the switch and log in as admin.
2. Issue the following command to enable the Track Changes feature:

```
trackchangesset 1
```

A message appears, verifying that the Track Changes feature is on:

```
switch:admin> trackchangesset 1
Committing configuration...done.
switch:admin>
```

The output from the Track Changes feature is dumped to the system message log for the switch. Use the `errdump` or `errshow` command to view the log.

Items in the system message log created from the Track Changes feature are labeled `TRCK`; for example:

```
2004/08/24-08:45:43, [TRCK-1001], 212,, INFO, ras007, Successful login by user
admin.
```

To display the status of the Track Changes feature:

1. Connect to the switch and log in as admin.
2. Issue the `trackchangesshow` command.

The status of the Track Changes feature is displayed as either `on` or `off`. The display specifies whether the Track Changes feature is configured to send SNMP traps:

```
switch:admin> trackchangesshow
Track Changes status: ON
Track Changes generate SNMP-TRAP: NO
switch:admin>
```

To view the switch status policy threshold values:

1. Connect to the switch and log in as admin.
2. Issue the `switchstatuspolicyshow` command at the command line.

Whenever there is a switch change, an error message is logged and an SNMP `connUnitStatusChange` trap is sent.

For the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32, the output is similar to the following:

```
switch:admin> switchstatuspolicyshow
The current overall switch status policy parameters:
      Down      Marginal
-----
PowerSupplies  2          1
Temperatures   2          1
      Fans    2          1
      Flash   0          1
MarginalPorts  5          2
      FaultyPorts 2          1
      MissingSFPs 2          1
switch:admin>
```

For the Core Switch 2/64 and the SAN Director 2/128, the output is similar to the following:

```
switch:admin> switchstatuspolicyshow
The current overall switch status policy parameters:
          Down      Marginal
-----
PowerSupplies  3          0
Temperatures   2          1
Fans           2          1
WWN            0          1
CP              0          1
Blade          0          1
Flash          0          1
MarginalPorts  2          1
FaultyPorts    2          1
MissingSFPs    0          0
switch:admin>
```

The policy parameter determines the number of failed or inoperable units for each contributor that triggers a status change in the switch.

Each parameter can be adjusted so that a specific threshold must be reached before that parameter changes the overall status of a switch to MARGINAL or DOWN. For example, if the FaultyPorts DOWN parameter is set to 3, the status of the switch changes if 3 ports fail. Only one policy parameter needs to pass the MARGINAL or DOWN threshold to change the overall status of the switch.

These parameters determine the status of a switch:

- Number of faulty ports
- Missing GBICs
- Power supply status
- Temperature in enclosure
- Fan speed
- Port status
- ISL status

For detailed information about setting policy parameters, refer to the *HP StorageWorks Fabric OS 4.x Fabric Watch user guide*.

To set the switch status policy threshold values:

1. Connect to the switch and log in as admin.
2. Issue the `switchstatuspolicyset` command.

First, the current switch status policy parameter values are displayed. Then, you are prompted to enter values for each DOWN and MARGINAL threshold parameter:

3. Verify the threshold settings you have configured for each parameter.
4. Issue the `switchstatuspolicyshow` command to view your current switch status policy configuration.



**NOTE:** By setting the DOWN and MARGINAL value for a parameter to 0, 0 that parameter is no longer used in setting the overall status for the switch.

For the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32, the following example shows the command as executed on a SAN Switch 2/32. The output is similar on the SAN Switch 2/8V, SAN Switch 2/16V, and SAN Switch 4/32:

```
switch:admin> switchstatuspolicyset
To change the overall switch status policy parameters
The current overall switch status policy parameters:
      Down      Marginal
-----
FaultyPorts  2          1
MissingSFPs  0          0
PowerSupplies 2          1
Temperatures 2          1
      Fans  2          1
PortStatus   0          0
ISLStatus    0          0
Note that the value, 0, for a parameter, means that it is
NOT used in the calculation.
** In addition, if the range of settable values in the prompt is (0..0),
** the policy parameter is NOT applicable to the switch.
** Simply hit the Return key.
The minimum number of
  FaultyPorts contributing to
      DOWN status: (0..32) [2] 3
  FaultyPorts contributing to
      MARGINAL status: (0..32) [1] 2
  MissingSFPs contributing to
      DOWN status: (0..32) [0]
  MissingSFPs contributing to
      MARGINAL status: (0..32) [0]
  Bad PowerSupplies contributing to
      DOWN status: (0..2) [2]
  Bad PowerSupplies contributing to
      MARGINAL status: (0..2) [1]
  Bad Temperatures contributing to
      DOWN status: (0..5) [2]
  Bad Temperatures contributing to
      MARGINAL status: (0..5) [1]
  Bad Fans contributing to
      DOWN status: (0..6) [2]
  Bad Fans contributing to
      MARGINAL status: (0..6) [1]
  Down PortStatus contributing to
      DOWN status: (0..32) [0]
  Down PortStatus contributing to
      MARGINAL status: (0..32) [0]
  down ISLStatus contributing to
      DOWN status: (0..32) [0]
  down ISLStatus contributing to
      MARGINAL status: (0..32) [0]
Policy parameter set has been changed
```

For the Core Switch 2/64 and the SAN Director 2/128, the command output includes parameters related to CP cards.

# 3 Configuring standard security features

This chapter provides information and procedures for standard Fabric OS security features. Standard Fabric OS features include account and password management. Additional security is available when secure mode is enabled. For information about licensed security features available in Secure Fabric OS, refer to the *HP StorageWorks Secure Fabric OS user guide*.

This chapter contains the following sections:

- [Ensuring network security](#), page 39
- [Configuring the telnet interface](#), page 40
- [Blocking listeners](#), page 41
- [Accessing switches and fabrics](#), page 42
- [Creating and maintaining user-defined accounts](#), page 43
- [Changing an account password](#), page 45
- [Setting up RADIUS AAA service](#), page 46
- [Configuring for the SSL protocol](#), page 54
- [Configuring for SNMP](#), page 60
- [Configuring secure file copy](#), page 69
- [Setting the boot PROM password](#), page 70
- [Recovering forgotten passwords](#), page 73

## Ensuring network security

To ensure security, Fabric OS supports secure shell (SSH) encrypted sessions. SSH encrypts all messages, including the client's transmission of password during login. The SSH package contains a daemon (sshd), which runs on the switch. The daemon supports a wide variety of encryption algorithms such as Blowfish-CBC and AES.



**NOTE:** To maintain a secure network, you should avoid using telnet or any other unprotected application when you are working on the switch. For example, if you use telnet to connect to a machine, then start an SSH or secure telnet session from that machine to the switch, the communication to the switch is in clear text, and therefore is not secure.

The FTP protocol is also not secure. When you use FTP to copy files to or from the switch, the contents are in clear text. This includes the remote FTP server's login and password. This limitation affects the following commands: `savecore`, `configupload`, `configdownload`, and `firmwaredownload`.

Commands that require a secure login channel must be issued from an original SSH session. If you start an SSH session, use the `login` command to start a nested SSH session, commands that require a secure channel are rejected.

Fabric OS v4.4.0 and later supports SSH protocol v2.0 (ssh2). For more information on SSH, see the SSH IETF web site: <http://www.ietf.org/ids.by.wg/secsh.html>

Refer to *SSH, The Secure Shell; The Definitive Guide*, By Daniel J. Barrett and Richard Silverman; Published by O'Reilly.

Fabric OS v4.4.0 comes with the SSH server preinstalled; however, you must select and install the SSH client. For information on installing and configuring the F-Secure SSH client, see the web site: <http://www.f-secure.com>

## Configuring the telnet interface

Telnet is enabled by default. To prevent users from passing clear text passwords over the network when they connect to the switch, you can disable the telnet interface.



**NOTE:** Before disabling the telnet interface, make sure that you have installed SSH, or some other secure means of establishing a connection with the switch.

To disable telnet:

1. Connect to the switch and log in as admin.

HP recommends that you connect through some other means than telnet; for example, through SSH.

2. Enter the following command:

```
configure telnetd
```

3. In response to the System Services prompt, enter **y**.

4. In response to the telnetd prompt, enter **off**.

The telnet interface is disabled. If you entered the command during a standard telnet session, the session terminates.

Example:

```
switch:admin> configure telnetd
Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.
Configure...
ssl attributes (yes, y, no, n): [no]
http attributes (yes, y, no, n): [no]
snmp attributes (yes, y, no, n): [no]
rpcd attributes (yes, y, no, n): [no]
cfgload attributes (yes, y, no, n): [no]

[31454]: Read 1 license entries for generation 1.
[31454]: Read 1 license records.
System services (yes, y, no, n): [no] y

rstatd (on, off): [off]
rusersd (on, off): [off]
telnetd (on, off): [on] off
```

To enable telnet:

1. Connect to the switch through a means other than telnet (for example, SSH) and log in as admin.
2. Issue the following command:  

```
configure telneta
```
3. In response to the System Services prompt, enter `y`.
4. In response to the `telneta` prompt, enter `on`.

The telnet interface is enabled.

---

## Blocking listeners

HP StorageWorks switches block Linux® subsystem listener applications that are not used to implement supported features and capabilities. [Table 5](#) lists the listener applications that HP StorageWorks switches either block or do not start.

**Table 5** Blocked listener applications

Listener application	Core Switch 2/64 and SAN Director 2/128	SAN switches 2/8V, 2/16V, 2/32, 4/32
chargen	Do not start	Do not start
echo	Do not start	Do not start
daytime	Do not start	Do not start
discard	Do not start	Do not start
ftp	Do not start	Do not start
rexec	Block with packet filter	Do not start
rsh	Block with packet filter	Do not start
rlogin	Block with packet filter	Do not start
time	Block with packet filter	Do not start
rstats	Do not start	Do not start
rusers	Do not start	Do not start

## Accessing switches and fabrics

Table 6 lists the defaults for accessing hosts, devices, switches, and zones.

**Table 6** Access details

Area	Description
Hosts	Any host can access the fabric by SNMP.
	Any host can telnet to any switch in the fabric.
	Any host can establish an HTTP connection to any switch in the fabric.
	Any host can establish an API connection to any switch in the fabric.
Devices	All devices can access the management server.
	Any device can connect to any FC port in the fabric.
Switch Access	Any switch can join the fabric.
	All switches in the fabric can be accessed through serial port.
Zoning	Node WWNs can be used for WWN-based zoning.

# Creating and maintaining user-defined accounts

In addition to the default administrative and user accounts, Fabric OS supports up to 15 user-defined accounts in each logical switch (domain). These accounts expand your ability to track account access and audit administrative activities.

User-defined accounts can be specified as either admin or user level. Admin-level accounts allow up to two simultaneous login sessions. User-level accounts allow up to four simultaneous login sessions. The total number of simultaneous login sessions allowed per logical switch is 15.

You can change passwords on user-defined accounts as described in “[Changing an account password](#)” on page 45.

If the Track Changes feature is enabled, the system keeps track of account names and login attempts. (See “[Tracking and controlling switch changes](#)” on page 35 for details on enabling the Track Changes feature.)

For large enterprises, Fabric OS also supports RADIUS services, as described in “[Setting up RADIUS AAA service](#)” on page 46.

The following procedures are for operations you can perform on user-defined accounts.

The default administrative account is called admin.



**NOTE:** If you are operating in secure mode, you can perform these operations only on the primary FCS switch.

## To display account information

1. Connect to the switch and log in as admin.
2. Issue one of the following commands:
  - `userConfig --show -a` to show all account information for a logical switch.
  - `userConfig --show -b` to show all backup account information for a logical switch.
  - `userConfig --show name` to show account information for the specified account name.

Accounts with the admin role can display information about all accounts on the logical switch. Accounts with the user role can display only information about themselves.

## To create a user-defined account

1. Connect to the switch and log in as admin.

- Issue the following command:

```
userConfig --add username -r rolename [-d description]
```

where:

<i>username</i>	Specifies the account name, which must begin with an alphabetic character. The name can be from 8 to 40 characters. It is case sensitive and can contain alphabetic and numeric characters, the dot (.), and the underscore (_). It must be different from all other account names on the logical switch.
<i>-r rolename</i>	Specifies the role: either admin or user in nonsecure mode; admin, user, or nonfcsadmin in secure mode.
<i>-d description</i>	Optionally, adds a description to the account. The description field can be up to 40 printable ASCII characters long. The following characters are not allowed: asterisk ('), quotation mark ("), exclamation point (!), semi-colon (;), and colon (:).

- In response to the prompt, enter a password for the account.

The password is not displayed when you enter it on the command line.

Accounts with the admin role can create accounts. Accounts with the user role cannot.

## To delete a user-defined account

- Connect to the switch and log in as admin.
- Issue the following command:

```
userConfig --delete username
```

where:

<i>username</i>	Specifies the account name. You cannot delete the default accounts. An account cannot delete itself. All active CLI sessions for the deleted account are logged out.
-----------------	--

- At the prompt for confirmation, enter **y**.

Accounts with the admin role can delete user-defined accounts on the logical switch. Accounts with the user role cannot.

## To change account parameters

- Connect to the switch and log in as admin.

**2.** Enter the following command:

```
userConfig --change username [-r rolename] [-d description] [-e yes | no]
```

where:

*username* An option that changes the account attribute for *username*. The account must already exist.

*-r rolename* An option that changes the role: either admin or user in nonsecure mode; admin, user, or nonfcsadmin in secure mode.

An account cannot change its own role.

You can only change the role name of a user-defined account with a lower level of authorization.

*-d description* An option: the account description. The description field can be up to 40 printable ASCII characters long. The following characters are not allowed: asterisk (\*), quotation mark ("), exclamation point (!), semi-colon (;), and colon (:).

You can only change the description of a user-defined account with a lower level of authorization.

*-e* Optionally, enter *yes* to enable the account or enter *no* to disable it. If you disable an account, all active CLI sessions for that account are logged out. You can enable or disable user-defined or default accounts.

Accounts with the admin role can change information for accounts that have lesser permissions. Accounts with the user role cannot.

## To recover user-defined accounts

If a backup account exists (in secure mode), you can recover it with the following command:

```
userConfig --recover
```

The following conditions apply to recovering user accounts:

- Only accounts with admin or higher roles can recover accounts.
- The attributes in the backup database replace the attributes in the current account database.
- An event is stored in the system message log indicating that accounts have been recovered.

---

## Changing an account password

At each level of account access, you can change passwords for that account and accounts that have lesser privileges.

If you log in to a user account, you can only change that account's password.

If you log in to an admin account, you can change admin and user passwords. You must provide the old password when the account being changed has the same or higher privileges than the current login account. For example, when logged in as admin, you need admin passwords to change passwords for admin accounts (except when you change the default user account password at login), but you do not need user passwords to change passwords for user accounts.

A new password must have at least one character different than the old password. The following rules also apply to passwords:

- You cannot change passwords using SNMP.

- Password prompting is disabled when security mode is enabled.
- Starting with Fabric OS v4.4.0, admin level accounts can use Web Tools to change passwords.
- Starting with Fabric OS v3.2.0, you cannot change default account names.
- For information on password behavior when you upgrade (or downgrade) firmware, see “[Effects of firmware changes on accounts and passwords](#)” on page 80.

To change the password for the current login account:

1. Connect to the switch and log in as either admin or user.
2. Issue the following command:

```
passwd
```

3. Enter the requested information at the prompts.

To change the password for a different account:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
passwd name
```

where *name* is the name of the account.

3. Enter the requested information at the prompts.

If the named account has lesser privileges than the current account, the old password is not required. If the named account has equal or higher privileges than the current account, you are prompted to enter the old password.

## Setting up RADIUS AAA service

Fabric OS v3.2 and v4.4 support Remote Authentication Dial-in User Service (RADIUS) authentication, authorization, and accounting (AAA). When it is configured for RADIUS, the switch becomes a RADIUS client. In this configuration, authentication records are stored in the RADIUS host server database.

The RADIUS service supports accounting request and response packets so that accounting records can be centralized on the RADIUS server. The login account name, assigned role, password, and time accounting records are stored on the RADIUS server for each user.

By default, RADIUS service is disabled, so AAA services default to the switch local database.

To enable RADIUS service, access the CLI through an SSH connection so that the shared secret is protected. Multiple login sessions can configure simultaneously, and the last session to apply a change leaves its configuration in effect. After a configuration is applied, it persists after a reboot or an HA failover.

The configuration is chassis-based, so it applies to all logical switches (domains) on the switch and replicates itself on a standby CP card, if one is present. It is saved in a configuration upload and applied in a configuration download.

Configure at least two RADIUS servers so that if one fails, the other assumes service. You can set the configuration with both RADIUS service and local authentication enabled so that if all RADIUS servers do not respond (because of power failure or network problems), the switch uses local authentication.

Consider the following effects of the use of RADIUS service on other Fabric OS features:

- When RADIUS service is enabled, all account passwords must be managed on the RADIUS server. The Fabric OS mechanisms for changing switch passwords remain functional; however, such changes affect only the involved switches locally. They do not propagate to the RADIUS server, nor do they affect any account on the RADIUS server.

When RADIUS is set up for a fabric that contains a mix of switches running v4.4.0 and v3.2.0 or earlier, the way a switch authenticates users depends on whether a RADIUS server is set up for that switch. For a switch with RADIUS support and configuration, authentication bypasses the local password database. For a switch without RADIUS support or configuration, authentication uses the switch's local account names and passwords.

- When Secure Fabric OS secure mode is enabled, the following items apply:
  - Account passwords are distributed among all switches in the same fabric. An account that resides on several switches has the same password on all of them. This model applies with RADIUS integration; however, such distribution affects only the switch's local password database.
  - There are separate admin and nonfcsadmin roles in secure mode. A nonfcsadmin account on a RADIUS server cannot access FCS switches, even if the account is properly authenticated.
  - If a nonfcsadmin account on a RADIUS server logs in to a switch in nonsecure mode, the switch treats the role like the admin role, and grants the access.
- The following items apply to Advanced Web Tools:
  - Advanced Web Tools client and server keep a session open after a user is authenticated. A password change on a switch invalidates an open session and requires the user to log in again. When integrated with RADIUS, a switch password change on the RADIUS server does not invalidate an existing open session, although a password change on the local switch does.
  - If you cannot log in because of a RADIUS server connection problem, Advanced Web Tools displays a message indicating server outage.
- The following items apply to API:
  - When an older version of the API host library authenticates against a switch with RADIUS support, the host performs the login. However, the old host library does not recognize the role returned from the switch, which can result in the host displaying an incorrect read or write attribute for an account. The switch library performs the permission check again for individual API function calls.
  - API provides functions for RADIUS configuration that share the behavior of the aaaConfig CLI command.
- The following items apply to both Advanced Web Tools and API:
  - Users can log in using account names and passwords configured on the RADIUS server and gain access with the switch roles defined there.
  - Users can log in through API using account names and passwords configured on the RADIUS server and gain access with the switch roles defined there.
  - When a proxy switch is used, the switch-side component performs authentication on the proxy switch, rather than on the destination switch. Therefore, to use RADIUS in this environment, you must configure on the proxy switch.

# Configuring the RADIUS server

You must know the switch IP address or name to connect to switches. Use the ipaddrshow command to display a switch IP address.

For the Core Switch 2/64 and the SAN Director 2/128 (chassis-based systems), the switch IP addresses are aliases of the physical Ethernet interfaces on the CP cards. When specifying client IP addresses for the logical switches in such systems, make sure that the CP card IP addresses are used. For accessing both the active and standby CP card, and for the purpose of HA failover, both of the CP card IP addresses should be included in the RADIUS server configuration.

User accounts should be set up by their true network-wide identity, rather than by the account names created on a Fabric OS switch. Along with each account name, the administrator should assign appropriate switch access roles. To manage a nonsecure fabric, these roles can be user or admin. To manage a secure fabric, these roles can be user, admin, or nonfcsadmin.

When they log in to a switch configured with RADIUS, users enter their assigned RADIUS account names and passwords at the prompt. After RADIUS server authenticates a user, it responds with the assigned switch role in HP Vendor-Specific Attribute (VSA) as defined in the RFC. An Authentication-Accept response without such VSA role assignment automatically assigns the user role.

The following sections describe how to configure a RADIUS server to support HP clients under different operating systems.

## Windows 2000

Use these procedures to add a client to the RADIUS server and create remote access policies for Fabric OS user and admin roles.

To add a RADIUS client:

1. From the Windows Start menu, select **Programs > Administrative Tools > Internet Authentication Service** to bring up the Internet Authentication Service window.
2. In the Internet Authentication Service window, right-click the **RADIUS Clients** folder and select **New RADIUS Client**.
3. In the New RADIUS Client window:
  - In the Friendly name space, enter a name for the switch that allows you to identify it easily.
  - In the Client Address (IP or DNS) space, enter the IP address of the switch.
4. Click **Next**.
5. In the next window, enter and confirm the shared secret, in the spaces provided. Make sure the shared secret matches that configured on the switch (as described in "[To add a RADIUS server to the switch configuration:](#)" on page 52).
6. Click **Finish**.

The new client friendly name appears in the list of clients. Should you need to change the shared secret, right-click the client, select **Properties**, and change the secret in the properties window.

To create user and admin remote access policies:

1. From the Windows Start menu, select **Programs > Administrative Tools > Internet Authentication Service** to bring up the Internet Authentication Service window.
2. If you do not already have Windows groups set up, use standard Windows procedures to set up a Windows group of login names assigned to the user role and another Windows group of login names assigned to the admin role.
3. Right-click the **Remote Access Policies** icon folder and select **New Remote Access Policy**.
4. In the New Remote Access Policy Wizard window, click **Next**.

5. In the Set Up a Custom Policy window:

  - a. Select the **Custom policy** radio button.
  - b. Enter a policy name for the user role (for example, **HP User**) in the space provided.
  - c. Click **Next**.
6. In the Select Attribute window, select **Windows-Groups** and click **Add**.
7. In the Select Groups window:

  - a. Enter the name of the Windows group that contains login names assigned to the user role.
  - b. Click **Check Names**.

When the system finds the Windows group, it underlines the name.
8. Click **OK**.
9. In the Group window, check that the Windows group is listed, and click **OK**.
10. In the Policy Conditions window, check that the policy name is listed (for example, **HP User**) and click **Next**.
11. In the Permissions window, select the **Grant remote access permission** radio button, and click **Next**.
12. In the Profile window, click **EDIT PROFILE**.
13. In the Edit Dial-in Profile window, click the **Authentication** tab.
14. In the Authentication tab:

  - Uncheck these check boxes:
    - Microsoft Encryption (MSCHAPv2)
    - Microsoft Encryption (MSCHAP)
  - Check these check boxes:
    - Encrypted Authentication (CHAP)
    - Unencrypted Authentication (PAP, SPAP)
15. Select the **Advanced** tab.
16. In the Advanced tab, click **Add**.
17. In the Add Attributes window, select **Vendor-specific** and click **Add**.
18. In the Multivalued Attribute Information window, click **ADD**.
19. In the Vendor-Specific Attribute Information window:

  - a. Select the **Enter Vendor Code** radio button and enter 1588 in the space provided.
  - b. Select the **Yes. It conforms.** radio button.
  - c. Click **Configure Attribute**.
20. In the Configure VSA (RFC Compliant) window, enter the following information in the spaces provided:

  - a. Vendor-Assigned Attribute Number: 1
  - b. Attribute Format: string
  - c. Attribute Value: user
21. Click **OK**.
22. Click **OK** or **Close** in each window until you reach the New Remote Access Policy Wizard.
23. Click **Next**.
24. Click **Finish**.

**25.** Repeat the procedure to set the admin remote access policy, with these differences:

- In [step 5](#), enter a policy name for the admin role (for example, HP Admin) in the space provided.
- In [step 7](#), enter the name of the Windows group that contains login names assigned to the admin role.
- In [step 20](#), enter admin in the Attribute Value space.

## Linux

Use the following procedure on a Linux FreeRADIUS server to:

- Set up a vendor dictionary file and include it in the system dictionary file.
- Identify a switch as a RADIUS client.
- Set up user accounts and roles.
- Test the configuration.

1. Log in to the server and change directory to the RADIUS configuration file directory. Typically, this directory is located at /usr/local/etc/raddb.
2. Use a text editor to create a vendor dictionary file called `dictionary.brocade` and enter the following lines into the file:

```
#  
# dictionary.brocade  
#  
VENDOR      Brocade          1588  
#  
# attributes  
#  
ATTRIBUTE  Brocade-Auth-Role 1      string      Brocade
```

**3.** Save `dictionary.brocade`.

**4.** Open the system dictionary file in a text editor and add this line:

```
$INCLUDE dictionary.brocade
```

The dictionary file is located in the RADIUS configuration directory.

**5.** Save the dictionary file.

**6.** Open the `client.config` file in a text editor and add the switches that are to be configured as RADIUS clients. For example, to configure the switch at IP address 10.32.170.59 as a client:

```
client 10.32.170.59  
secret      = Secret  
shortname   = Testing Switch  
nastype     = other
```

The `client.config` file is located in the RADIUS configuration directory.

In this example, the switch name is Testing Switch and its shared secret is Secret. Make sure that the shared secret matches that configured on the switch (see "[To add a RADIUS server to the switch configuration:](#)" on page 52).

**7.** Save `client.config`.

**8.** Open the `user` file in a text editor and add user names and roles for users who will be accessing the switch. For example, to set up an account called `JohnDoe` with the admin role:

```
JohnDoe Auth-Type := Local, User-Password == "johnPassword"  
HP-Auth-Role = "admin"
```

The `user` file is located in the RADIUS configuration directory.

**9.** Save the `user` file.

**10.** Enter this command to start the RADIUS server:

```
/usr/local/sbin/radiusd
```

**11.** Log in to a client switch and use the `aaaconfig` command to configure it as a client and enable RADIUS service, as described in "[To add a RADIUS server to the switch configuration:](#)" on page 52 and "[To enable or disable RADIUS service:](#)" on page 52.

**12.** Log out.

When you log in to the switch again, RADIUS service is in force.

## Configuring the switch

The following procedures show how to use the `aaaconfig` command to set up a switch for RADIUS service.

To display the current RADIUS configuration

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> aaaConfig --show
```

If a configuration exists, its parameters are displayed. If RADIUS service is not configured, only the parameter heading line is displayed. Parameters include:

Position	The order in which servers are contacted to provide service
Server	The server names or IP addresses
Port	The server ports
Secret	The shared secrets
Timeouts	The length of time servers have to respond before the next server is contacted
Authentication	The type of authentication being used on servers

To add a RADIUS server to the switch configuration:

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> aaaConfig --add server [-p port] [-s secret] [-t timeout]
[-a pap | chap]
```

where:

<i>server</i>	Enter either a server name or IP address. Avoid duplicating server listings (that is, listing the same server once by name and again by IP address). Up to five servers can be added to the configuration.
<i>-p port</i>	Optional: enter a server port. The default is port 1812.
<i>-s secret</i>	Optional: enter a shared secret. The default is sharedsecret. Secrets can be from 8 to 40 alphanumeric characters long. Make sure that the secret matches that configured on the server.
<i>-t timeout</i>	Optional: enter the time (in seconds) the server has to respond before the next server is contacted. The default is three seconds. Timeout values can range from 1 to 30 seconds.
<i>-a</i>	Optional: specify that the PAP protocol be used instead of the CHAP protocol for packets traveling between the switch and the server.

To enable or disable RADIUS service:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
switch:admin> aaaConfig --radius on | off
```

Specifying **on** enables the service; specifying **off** disables it.

At least one RADIUS server must be configured before you can enable RADIUS service.

If no RADIUS configuration exists, turning it on triggers an error message. When the command succeeds, the event log indicates that the configuration is enabled or disabled.

To delete a RADIUS server from the configuration:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
switch:admin> aaaConfig --remove server | all
```

where:

<i>server</i>	Servers are listed by either name or IP address. Enter either the name or IP address of the server to be removed.
<i>all</i>	Enter this keyword to remove all servers. If RADIUS service is enabled, this removes all but the server in the first position. If RADIUS service is disabled, all servers are removed.

3. At the prompt, enter **y** to complete the command.

When the command succeeds, the event log indicates that the server is removed.

To change a RADIUS server configuration:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
switch:admin> aaaConfig --change server [-p port] [-s secret] [-t  
timeout] [-a pap | chap]
```

where:

<i>server</i>	Servers are listed by either name or IP address. Enter either the name or IP address of the server to be changed.
<i>-p port</i>	Optional: enter a server port.
<i>-s secret</i>	Optional: enter a shared secret.
<i>-t timeout</i>	Optional: enter the length of time (in seconds) the server has to respond before the next server is contacted.
<i>-a pap   chap</i>	Optional: specify that the PAP protocol be used instead of the CHAP protocol for packets traveling between the switch and the server.

To change the order in which RADIUS servers are contacted for service:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
switch:admin> aaaConfig --move server to_position
```

where:

<i>server</i>	Servers are listed by either name or IP address. Enter either the name or IP address of the server whose position is to be changed.
<i>to_position</i>	Enter the position number to which the server is to be moved.

When the command succeeds, the event log indicates that a server configuration changed.

## Enabling and disabling local authentication

It is useful to enable local authentication so that the switch can take over authentication locally if the RADIUS servers fail to respond because of power outage or network problems. To enable or disable local authentication, issue the following command:

```
switch:admin> aaaConfig --switchdb on | off
```

Specifying **on** enables local authentication; specifying **off** disables it.

When local authentication is enabled and RADIUS servers fail to respond, you can log in to the default switch accounts (admin and user) or any user-defined account. You must know the passwords of these accounts.

RADIUS authentication must be enabled when local database authentication is turned off from the **on** state; otherwise, an error is returned.

Because local database authentication may be automatically disabled or enabled when enabling or disabling RADIUS authentication, you should set the local database authentication explicitly to enabled or disabled after setting the desired RADIUS authentication configuration.

When the command succeeds, the event log indicates that local database authentication is disabled or enabled.

# Configuring for the SSL protocol

Fabric OS v4.4.0 and later supports secure sockets layer (SSL) protocol, which provides secure access to a fabric through web-based management tools like Advanced Web Tools. SSL support is a standard Fabric OS feature; it is independent of Secure Fabric OS, which requires a license and separate certification.

Switches configured for SSL grant access to management tools through hypertext transfer protocol-secure links (which begin with `https://`) instead of standard links (which begin with `http://`).

SSL uses public key infrastructure (PKI) encryption to protect data transferred over SSL connections. PKI is based on digital certificates obtained from an Internet Certificate Authority (CA), which acts as the trusted key agent.

Certificates are based on the switch IP address or fully qualified domain name (FQDN), depending on the issuing CA. If you change a switch IP address or FQDN after activating an associated certificate, you might have to obtain and install a new certificate. Check with the CA to verify this possibility, and plan these types of changes accordingly.

## Browser and Java™ support

Fabric OS supports the following Web browsers for SSL connections:

- Internet Explorer (Microsoft Windows)
- Mozilla (Solaris and Redhat Linux)

In countries that allow the use of 128-bit encryption, you should use the latest version of your browser. For example, Internet Explorer 6.0 and later supports 128-bit encryption by default. You can display the encryption support (called "cipher strength") using the Internet Explorer Help:About menu option. If you are running an earlier version of Internet Explorer, you might be able to download an encryption patch from the Microsoft Web site at <http://www.microsoft.com>.

You should upgrade to the Java 1.4.2\_03 Plug-in on your management workstation. To find the Java version that is currently running, open the Java console and look at the first line of the window.

For more details on levels of browser and Java support, refer to the *HP StorageWorks Fabric OS 4.x Advanced Web Tools user guide*.

## Summary of SSL procedures

You configure for SSL by obtaining, installing, and activating digital certificates for SSL support. Certificates are required on all switches that are to be accessed through SSL.

You also need to install a certificate to the Java Plug-in on the management workstation, and you might need to add a certificate to your Web browser.

Configuring for SSL involves these major steps, which are shown in detail in the next sections:

1. Choose a CA.
2. On each switch:
  - a. Generate a public/private key (`seccertutil genkey` command).
  - b. Generate a certificate signing request (CSR) (`seccertutil gencsr` command) and store the CSR on an FTP server (`seccertutil export` command).
3. Obtain the certificates from the CA.

You can request a certificate from a CA through a Web browser. After you request a certificate, the CA either sends certificate files by e-mail (public) or gives access to them on a remote host (private). Typically, the CA provides the certificate files listed in [Table 7](#).

**Table 7** SSL certificate files

Certificate file	Description
name.crt	The switch certificate.
nameRoot.crt	The root certificate. Typically, this certificate is already installed in the browser, but if not, you must install it.
nameCA.crt	The CA certificate. It is not necessary to install this, but you can if you want the CA name to be displayed in the browser window.

4. On each switch:
  - a. Install the certificate.
  - b. Activate the certificate.
5. If necessary, install the root certificate to the browser on the management workstation.
6. Add the root certificate to the Java Plug-in keystore on the management workstation.

## Choosing a CA

To ease maintenance and allow secure out-of-band communication between switches, consider using one CA to sign all management certificates for a fabric. If you use different CAs, management services operate correctly, but the Web Tools Fabric Events button is unable to retrieve events for the entire fabric.

[Table 8](#) lists recommended Certificate Authorities. Each CA has slightly different requirements; for example, some generate certificates based on IP address, while others require an FQDN, and most require a 1024-bit public/private key while some might accept a 2048-bit key. Consider your fabric configuration, check CA Web sites for requirements, and gather all the information that the CA requires.

**Table 8** Recommended CAs

Certificate authority	Web Site
Verisign	<a href="http://www.verisign.com">www.verisign.com</a>
Entrust	<a href="http://www.entrust.com">www.entrust.com</a>
InstantSSL	<a href="http://www.instantssl.com">www.instantssl.com</a>
GeoTrust	<a href="http://www.geotrust.com">www.geotrust.com</a>

## Generating a public/private key

Perform this procedure on each switch:

1. Connect to the switch and log in as admin.
2. Issue the following command to generate a public/private key pair:

```
switch:admin> seccertutil genkey
```

The system reports that this process disables secure protocols, delete any existing CSR, and delete any existing certificates.

3. Respond to the prompts to continue and select the key size:

```
Continue (yes, y, no, n): [no] y
Select key size [1024 or 2048]: 1024
Generating new rsa public/private key pair
Done.
```

Because CA support for the 2048-bit key size is limited, you should select 1024 in most cases.

## Generating and storing a CSR

After generating a public/private key (see “[Generating a public/private key](#)” on page 55 earlier), perform this procedure on each switch:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
switch:admin> seccertutil gencsr
```

3. Enter the requested information:

```
Country Name (2 letter code, eg, US):US
State or Province Name (full name, eg, California):California
Locality Name (eg, city name):San Jose
Organization Name (eg, company name):HP
Organizational Unit Name (eg, department name):Eng
Common Name (Fully qualified Domain Name, or IP address): 192.1.2.3
Generating CSR, file name is: 192.1.2.3.csr
Done.
```

Your CA might require specific codes for Country, State or Province, Locality, Organization, and Organizational Unit names. Make sure that your spelling is correct and matches the CA requirements. If the CA requires that the Common Name be specified as an FQDN, make sure that the fully qualified domain name is set on the domain Name Server.

4. Issue the following command to store the CSR:

```
switch:admin> seccertutil export
```

5. Enter the requested information:

```
Select protocol [ftp or scp]: ftp
Enter IP address: 192.1.2.3
Enter remote directory: path_to_remote_directory
Enter Login Name: your account
Enter Password: your password
Success: exported CSR.
```

If you are set up for secure file copy protocol, you can select it; otherwise, select ftp. Enter the IP address of the switch on which you generated the CSR. Enter the remote directory name of the FTP server to which the CSR is to be sent. Enter your account name and password on the server.

## Obtaining certificates

Check the instructions on the CA web site; then, perform this procedure for each switch:

1. Generate and store the CSR as described in “[Generating and storing a CSR](#)” on page 56.
2. Open a Web browser window on the management workstation and go to the CA web site. Follow the instructions to request a certificate. Locate the area in the request form that is provided for you to paste the CSR.
3. Through a telnet window, connect to the switch and log in as admin.
4. Issue the following command:

```
switch:admin> seccertutil showcsr
```

The contents of the CSR is displayed.

5. Locate the section that begins with BEGIN CERTIFICATE REQUEST and ends with END CERTIFICATE REQUEST.
6. Copy and paste this section (including the BEGIN and END lines) into the area provided in the request form; then, follow the instructions to complete and send the request.

It may take several days to receive the certificates. If the certificates arrive by e-mail, save them to an FTP server. If the CA provides access to the certificates on an FTP server, make note of the path name and make sure you have a login name and password on the server.

## Installing a switch certificate

Perform this procedure on each switch:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
switch:admin> seccertutil import
```

3. Select a protocol, enter the IP address of the host on which the switch certificate is saved, and enter your login name and password:

```
Select protocol [ftp or scp]: ftp
Enter IP address: 192.10.11.12
Enter remote directory: path_to_remote_directory
Enter certificate name (must have ".crt" suffix): 192.1.2.3.crt
Enter Login Name: your_account
Enter Password: *****
Success: imported certificate [192.1.2.3.crt].
To use this certificate, run the configure command to activate it
```

The certificate downloads to the switch.

## Activating a switch certificate

Enter the `configure` command and respond to the prompts that apply to SSL certificates:

SSL attributes	Enter <code>yes</code> .
Certificate File	Enter the name of the switch certificate file: for example, <code>192.1.2.3.crt</code> .
CA Certificate File	If you want the CA name to be displayed in the browser window, enter the name of the CA certificate file; otherwise, skip this prompt.
Select length of crypto key	Enter the encryption key length (40, 56, or 128).
HTTP attributes	Enter <code>yes</code> .
Secure HTTP enabled	Enter <code>yes</code> .

### Example

```
Configure...
System services (yes, y, no, n): [no]
    ssl attributes (yes, y, no, n): [no] yes
Certificate File. (filename or none): [10.33.13.182.crt] 192.1.2.3.crt
    CA Certificate File. (filename or none): [none]
    Select length of crypto key.
        (Valid values are 40, 56, and 128.): (40..128) [128]
http attributes (yes, y, no, n): [no] yes
HTTP Enabled (yes, y, no, n): [yes] no
    Secure HTTP Enabled (yes, y, no, n): [no] yes
```

After you exit the `configure` command, the HTTP daemon restarts automatically to handle HTTPS requests.

## Configuring the browser

The root certificate might already be installed on your browser, but if not, you must install it. To see whether it is already installed, check the certificate store on your browser.

The next procedures are guides for installing root certificates to Internet Explorer and Mozilla browsers. For more detailed instructions, refer to the documentation that came with the certificate.

To check and install root certificates on Internet Explorer:

1. From the browser Tools menu, select **Internet Options**.
2. Select the **Content** tab.
3. Click **Certificates**.
4. Click the various tabs and scroll the lists to see if the root certificate is listed. If it is listed, you do not need to install it.
5. If the certificate is not listed, click **Import**.
6. Follow the instructions in the Certificate Import wizard to import the certificate.

To check and install root certificates on Mozilla:

1. From the browser Edit menu, select **Preferences**.
2. In the left pane of the Preferences window, expand the **Privacy & Security** list and select **Certificates**.
3. In the right pane, click **Manage Certificates**.
4. In the next window, click the **Authorities** tab.

5. Scroll the authorities list to see if the root certificate is listed. (For example, its name may have the form nameRoot.crt.) If it is listed, you do not need to install it; forgo the remainder of this procedure.
6. If the certificate is not listed, click **Import**.
7. Browse to the certificate location and select the certificate. (For example, select *nameRoot.crt*.)
8. Click **Open** and follow the instructions to import the certificate.

## Installing a root certificate to the Java Plug-in

For information on Java requirements, see "Browser and Javatm support" on page 54.

This procedure is a guide for installing a root certificate to the Java Plug-in on the management workstation. If the root certificate is not already installed to the plug-in, you should install it. For more detailed instructions, refer to the documentation that came with the certificate and to the Sun Microsystems Web site ([www.sun.com](http://www.sun.com)).

1. Copy the root certificate file from its location on the FTP server to the Java Plug-in bin. For example, the bin location might be:

```
C: \program files\java\j2re1.4.2_03\bin
```

2. Open a Command Prompt window and change directory to the Java Plug-in bin.
3. Issue the keytool command and respond to the prompts:

```
C:\Program Files\Java\j2re1.4.2_03\bin> keytool -import -alias RootCert -file RootCert.crt -keystore ..\lib\security\RootCerts
Enter keystore password: changeit
Owner: CN=HP, OU=Software, O=HP Communications, L=San Jose, ST=California, C=US
Issuer: CN=HP, OU=Software, O=HP Communications, L=San Jose, ST=California, C=US
Serial number: 0
Valid from: Thu Jan 15 16:27:03 PST 2004 until: Sat Feb 14 16:27:03 PST 2004
Certificate fingerprints:
      MD5: 71:E9:27:44:01:30:48:CC:09:4D:11:80:9D:DE:A5:E3
      SHA1: 06:46:C5:A5:C8:6C:93:9C:FE:6A:C0:EC:66:E9:51:C2:DB:E6:4F:A1
Trust this certificate? [no]: yes
Certificate was added to keystore
```

In the example, *changeit* is the default password and *RootCert* is an example root certificate name.

## Displaying and deleting certificates

**Table 9** summarizes the commands for displaying and deleting certificates. For details on the commands, refer to the *HP StorageWorks Fabric OS 4.x command reference guide*.

**Table 9** Commands for displaying and deleting SSL certificates

Command	Description
seccertutil show	Displays the state of the SSL key and a list of installed certificates.
seccertutil show <i>filename</i>	Displays the contents of a specific certificate.
seccertutil showcsr	Displays the contents of a CSR.
seccertutil delete <i>filename</i>	Deletes a specified certificate.
seccertutil delcsr	Deletes a CSR.

## Troubleshooting certificates

If you receive messages in the browser or in a pop-up window when logging in to the target switch using HTTPS, see [Table 10](#).

**Table 10** SSL messages and actions

Message	Action
The page cannot be displayed	The SSL certificate is not installed correctly or HTTPS is not enabled correctly. Make sure that the certificate has not expired, that HTTPS is enabled, and that certificate file names are configured correctly.
The security certificate was issued by a company you have not chosen to trust....	The certificate is not installed in the browser. Install it as described in " <a href="#">Configuring the browser</a> " on page 58.
The security certificate has expired or is not yet valid	Either the certificate file is corrupted or it needs to be updated. Click <b>View Certificate</b> to verify the certificate content. If it is corrupted or out of date, obtain and install a new certificate.
The name on the security certificate is invalid or does not match the name of the site file	The certificate is not installed correctly in the Java Plug-in. Install it as described in " <a href="#">Installing a root certificate to the Java Plug-in</a> " on page 59.
This page contains both secure and nonsecure items. Do you want to display the nonsecure items?	Click <b>No</b> in this pop-up window. The session opens with a closed lock on the lower-right corner of the browser, indicating an encrypted connection.

## Configuring for SNMP

You can configure for the automatic transmission of Simple Network Management Protocol (SNMP) information to management stations. SNMPv3 and SNMPv1 are supported.

The configuration process involves configuring the SNMP agent and configuring SNMP traps. The following commands are used in the process:

- Use the `configure` command to set the security level. You can specify no security, authentication only, or authentication and privacy.
- Use the `snmpconfig` command to configure the SNMP agent and traps for SNMPv3 or SNMPv1 configurations.
- If necessary for backward compatibility, you can use these legacy commands to configure for SNMP v1:
  - Use the `agtcfgshow`, `agtcfgset`, and `agtcfgdefault` commands to configure the SNMPv1 agent.
  - Use the `snmpmibcapset` command to filter at the trap level and the `snmpmibcapshow` command to display the trap filter values.

Associated with the HP-specific StorageWorks MIB (SW-MIB), this Management Information Base (MIB) monitors HP StorageWorks switches specifically.

- Fibre Alliance MIB trap

Associated with the Fibre Alliance MIB (FA-MIB), this MIB manages SAN switches and devices from any company that complies with Fibre Alliance specifications.

If you use both SW-MIB and FA-MIB, you may receive duplicate information. You can disable the FA-MIB, but the SW-MIB cannot be disabled.

You can also use these additional MIBs and their associated traps:

- HA-MIB (for the Core Switch 2/64 and SAN Director 2/128)
- SW-EXTTRAP includes the swSsn (Software Serial Number) as a part of HP SW traps. It is also used with the legacy SAN Switched Integrated/64 to provide detailed group information for a particular trap.

For more information on HP support for SNMP, refer to the *HP StorageWorks Fabric OS 4.x features overview guide*.

For information on HP MIBs, refer to the *HP StorageWorks Fabric OS 4.x Management Information Base reference guide*.

For information on the specific commands used in these procedures, refer to online help or to the *HP StorageWorks Fabric OS 4.x command reference guide*.

## Setting the security level

Use the `configure` command to set the security level (called SNMP attributes). You can specify no security, authentication only, or authentication and privacy. For example, to configure for authentication and privacy:

```
switch:admin> configure

Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.

Configure...

System services (yes, y, no, n): [no]
ssl attributes (yes, y, no, n): [no]
http attributes (yes, y, no, n): [no]
snmp attributes (yes, y, no, n): [no] y

Select SNMP Security Level:
(0 = No security, 1 = Authentication only, 2 = Authentication
and Privacy): (0..2) [0] 2
```

## Using the `snmpconfig` command

Use the `snmpconfig --set` command to change either the SNMPv3 or SNMPv1 configuration. You can also change access control, MIB capability, and system group.

To change the SNMPv3 configuration, use the following as an example:

```
switch:admin> snmpconfig --set snmpv3

SNMPv3 user configuration:
User (rw): [snmpadmin1] adminuser
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3] 1
New Auth Passwd:
Verify Auth Passwd:
Priv Protocol [DES(1)/noPriv[2]): (1..2) [2] 1
New Priv Passwd:
Verify Priv Passwd:
User (rw): [snmpadmin2] shauser
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3] 2
New Auth Passwd:
Verify Auth Passwd:
Priv Protocol [DES(1)/noPriv[2]): (1..2) [2] 1
New Priv Passwd:
Verify Priv Passwd:
User (rw): [snmpadmin3] nosec
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv[2]): (2..2) [2]
User (ro): [snmpuser1]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (3..3) [3]
Priv Protocol [DES(1)/noPriv[2]): (2..2) [2]
User (ro): [snmpuser2]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (3..3) [3]
Priv Protocol [DES(1)/noPriv[2]): (2..2) [2]
User (ro): [snmpuser3]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (3..3) [3]
Priv Protocol [DES(1)/noPriv[2]): (2..2) [2]

SNMPv3 trap recipient configuration:
Trap Recipient's IP address in dot notation: [0.0.0.0]
192.168.45.90
UserIndex: (1..6) [1]
Trap recipient Severity level : (0..5) [0] 4
Trap Recipient's IP address in dot notation: [0.0.0.0]
192.168.45.92
UserIndex: (1..6) [2]
Trap recipient Severity level : (0..5) [0] 2
Trap Recipient's IP address in dot notation: [0.0.0.0]
Committing configuration...done.
```

To change the SNMPv1 configuration, use the following as an example:

```
switch:admin> snmpconfig --set snmpv1

SNMP community and trap recipient configuration:
Community (rw): [Secret C0de] admin
Trap Recipient's IP address in dot notation: [0.0.0.0] 10.32.225.1
Trap recipient Severity level : (0..5) [0] 1
Community (rw): [OrigEquipMfr]
Trap Recipient's IP address in dot notation: [10.32.225.2]
Trap recipient Severity level : (0..5) [1]
Community (rw): [private]
Trap Recipient's IP address in dot notation: [10.32.225.3]
Trap recipient Severity level : (0..5) [2]
Community (ro): [public]
Trap Recipient's IP address in dot notation: [10.32.225.4]
Trap recipient Severity level : (0..5) [3]
Community (ro): [common]
Trap Recipient's IP address in dot notation: [10.32.225.5]
Trap recipient Severity level : (0..5) [4]
Community (ro): [FibreChannel]
Trap Recipient's IP address in dot notation: [10.32.225.6]
Trap recipient Severity level : (0..5) [5]
Committing configuration...done.
```

To change the accessControl configuration, use the following as an example:

```
switch:admin> snmpconfig --set accessControl

SNMP access list configuration:
Access host subnet area in dot notation: [0.0.0.0] 192.168.0.0
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0] 10.32.148.0
Read/Write? (true, t, false, f): [true] f
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0] 10.33.0.0
Read/Write? (true, t, false, f): [true] f
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Committing configuration...done.
```

To display the mibCapability configuration, use the following as an example:

```
switch:admin> snmpconfig --show mibCapability
FA-MIB: YES
FICON-MIB: YES
HA-MIB: YES
SW-TRAP: YES
    swFCPortScn: YES
    swEventTrap: YES
    swFabricWatchTrap: YES
    swTrackChangesTrap: NO
FA-TRAP: YES
    connUnitStatusChange: YES
    connUnitEventTrap: NO
    connUnitSensorStatusChange: YES
    connUnitPortStatusChange: YES
SW-EXTTRAP: NO
FICON-TRAP: NO
HA-TRAP: YES
    fruStatusChanged: YES
    cpStatusChanged: YES
    fruHistoryTrap: NO
```

To change the systemGroup configuration to default, use the following as an example:

```
switch:admin> snmpconfig --default systemGroup
*****
This command will reset the agent's system group configuration back
to factory default
*****
sysDescr = Fibre Channel Switch
sysLocation = End User Premise
sysContact = Field Support
authTraps = 0 (OFF)

*****
Are you sure? (yes, y, no, n): [no] y
```

## Using legacy commands for SNMPv1

Use the `snmpconfig` command to configure the SNMPv1 agent and traps (see “[Using the snmpconfig command](#)” on page 61). However, if necessary for backward compatibility, you can choose to use legacy commands.

Use the agtcfgshow command to display SNMP agent configuration information. For example:

```
switch:admin> agtcfgshow
Current SNMP Agent Configuration
    Customizable MIB-II system variables:
        sysDescr = FC Switch
        sysLocation = End User Premise
        sysContact = Field Support.
        authTraps = 1 (ON)

SNMPv1 community and trap recipient configuration:
    Community 1: Secret C0de (rw)
        Trap recipient: 192.168.1.51
        Trap recipient Severity level: 4
    Community 2: OrigEquipMfr (rw)
        Trap recipient: 192.168.1.26
        Trap recipient Severity level: 0
    Community 3: private (rw)
        No trap recipient configured yet
    Community 4: public (ro)
        No trap recipient configured yet
    Community 5: common (ro)
        No trap recipient configured yet
    Community 6: FibreChannel (ro)
        No trap recipient configured yet

SNMP access list configuration:
Entry 0: Access host subnet area 192.168.64.0 (rw) ]
Entry 1: No access host configured yet
Entry 2: No access host configured yet
Entry 3: No access host configured yet
Entry 4: No access host configured yet
Entry 5: No access host configured yet
```

Use the agtcfgset command to modify the SNMP configuration values. For example:

```
switch:admin> agtcfgset

        Customizing MIB-II system variables ...

At each prompt, do one of the followings:
o <Return> to accept current value,
o enter the appropriate new value,
o <Control-D> to skip the rest of configuration, or
o <Control-C> to cancel any change.

To correct any input mistake:
<Backspace> erases the previous character,
<Control-U> erases the whole line,
sysDescr: [FC Switch]
sysLocation: [End User Premise]
sysContact: [Field Support.]
authTrapsEnabled (true, t, false, f): [true]

SNMP community and trap recipient configuration:
Community (rw): [Secret C0de]
Trap Recipient's IP address in dot notation: [192.168.1.51]
Trap recipient Severity level : (0..5) [0] 3
Community (rw): [OrigEquipMfr]
Trap Recipient's IP address in dot notation: [192.168.1.26]
Trap recipient Severity level : (0..5) [0]
Community (rw): [private]
Trap Recipient's IP address in dot notation: [0.0.0.0] 192.168.64.88
Trap recipient Severity level : (0..5) [0] 1
Community (ro): [public]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Community (ro): [common]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Community (ro): [FibreChannel]
Trap Recipient's IP address in dot notation: [0.0.0.0]

SNMP access list configuration:
Access host subnet area in dot notation: [0.0.0.0] 192.168.64.0
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Committing configuration...done.
value = 1 = 0x1
```

Use the **agtcfgdefault** command to reset the SNMP agent configuration to default values. For example:

```
switch:admin> agtcfgdefault
*****
This command will reset the agent's configuration back to factory default
*****
    Current SNMP Agent Configuration
    Customizable MIB-II system variables:
        sysDescr = Fibre Channel Switch.
        sysLocation = End User Premise
        sysContact = sweng
        authTraps = 0 (OFF)
SNMPv1 community and trap recipient configuration:
    Community 1: Secret C0de (rw)
        Trap recipient: 192.168.15.41
        Trap recipient Severity level: 4
    Community 2: OrigEquipMfr (rw)
        No trap recipient configured yet
    Community 3: private (rw)
        No trap recipient configured yet
    Community 4: public (ro)
        No trap recipient configured yet
    Community 5: common (ro)
        No trap recipient configured yet
    Community 6: FibreChannel (ro)
        No trap recipient configured yet
SNMP access list configuration:
    Entry 0: Access host subnet area 192.168.64.0 (rw) ]
    Entry 1: No access host configured yet
    Entry 2: No access host configured yet
    Entry 3: No access host configured yet
    Entry 4: No access host configured yet
    Entry 5: No access host configured yet
*****
Are you sure? (yes, y, no, n): [no] y
Committing configuration...done.
agent configuration reset to factory default
Current SNMP Agent Configuration
Customizable MIB-II system variables:
    sysDescr = Fibre Channel Switch.
    sysLocation = End User Premise
    sysContact = Field Support.
    authTraps = 0 (OFF)
SNMPv1 community and trap recipient configuration:
    Community 1: Secret C0de (rw)
        No trap recipient configured yet
    Community 2: OrigEquipMfr (rw)
        No trap recipient configured yet
    Community 3: private (rw)
        No trap recipient configured yet
    Community 4: public (ro)
        No trap recipient configured yet
    Community 5: common (ro)
        No trap recipient configured yet
    Community 6: FibreChannel (ro)
        No trap recipient configured yet
(output truncated)
```

Use the **snmpmibcapset** command to modify the options for configuring SNMP MIB traps. For example:

```
switch:admin> snmpmibcapset
The SNMP Mib/Trap Capability has been set to support
FE-MIB
SW-MIB
FA-MIB
FA-TRAP
FA-MIB (yes, y, no, n): [yes]
FICON-MIB (yes, y, no, n): [no] y
HA-MIB (yes, y, no, n): [no] y
SW-TRAP (yes, y, no, n): [no] y
    swFCPortScn (yes, y, no, n): [no]
    swEventTrap (yes, y, no, n): [no]
    swFabricWatchTrap (yes, y, no, n): [no]
    swTrackChangesTrap (yes, y, no, n): [no]
FA-TRAP (yes, y, no, n): [yes]
    connUnitStatusChange (yes, y, no, n): [no]
    connUnitEventTrap (yes, y, no, n): [no]
    connUnitSensorStatusChange (yes, y, no, n): [no]
    connUnitPortStatusChange (yes, y, no, n): [no]
SW-EXTTRAP (yes, y, no, n): [no] y
FICON-TRAP (yes, y, no, n): [no] y
    linkRNIDDeviceRegistration (yes, y, no, n): [no]
    linkRNIDDeviceDeRegistration (yes, y, no, n): [no]
    linkLIRRLListenerAdded (yes, y, no, n): [no]
    linkLIRRLListenerRemoved (yes, y, no, n): [no]
    linkRLIRFailureIncident (yes, y, no, n): [no]
HA-TRAP (yes, y, no, n): [no] y
    fruStatusChanged (yes, y, no, n): [no]
    cpStatusChanged (yes, y, no, n): [no]
    fruHistoryTrap (yes, y, no, n): [no]
Avoid-Duplicate-TRAP (yes, y, no, n): [no] y
switch:admin>
```

These notes apply to **snmpmibcapset** parameters for the FA-TRAP:

- **connUnitStatusChange** indicates that the overall status of the connectivity unit has changed. Its variables are:
  - **connUnitStatus** is the status of the connection unit
  - **connUnitState** is the state of the connection unit
- **connUnitEventTrap** indicates that the connectivity unit has generated an event. Its variables are:
  - **connUnitEventId** is the internal event ID
  - **connUnitEventType** is the type of this event
- **connUnitEventObject** is used with the **connUnitEventType** to identify the object to which the event refers.
- **connUnitEventDescr** is the description of the event.
- **connUnitSensorStatusChange** indicates that the status of the sensor associated with the connectivity unit has changed.
  - **connUnitSensorStatus** is the status indicated by the sensor.

- `connUnitPortStatusChange` indicates that the status of the sensor associated with the connectivity unit has changed.
- `connUnitPortStatus` shows overall protocol status for the port.
- `connUnitPortState` shows the user-specified state of the port hardware.

Use the `snmpmibcapshow` command to view the SNMP MIB trap setup. For example:

```
switch:admin> snmpmibcapshow
FA-MIB: YES
FICON-MIB: YES
HA-MIB: YES
SW-TRAP: YES
    swFCPortScn: YES
    swEventTrap: YES
    swFabricWatchTrap: YES
    swTrackChangesTrap: YES
FA-TRAP: YES
SW-EXTTRAP: YES
HA-TRAP: YES
    fruStatusChanged: YES
    cpStatusChanged: YES
    fruHistoryTrap: YES
```

## Configuring secure file copy

Use the `configure` command to specify that secure file copy (`scp`) be used for configuration uploads and downloads. For example:

```
switch:admin> configure

Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.

Configure...

System services (yes, y, no, n): [no] n
ssl attributes (yes, y, no, n): [no] n
http attributes (yes, y, no, n): [no] n
snmp attributes (yes, y, no, n): [no] n
rpcd attributes (yes, y, no, n): [no] n
cfgload attributes (yes, y, no, n): [no] y

Enforce secure config Upload/Download (yes, y, no, n): [no] y
switch:admin>
```

# Setting the boot PROM password

The boot PROM password provides an additional layer of security by protecting the boot PROM from unauthorized use. Setting a recovery string for the boot PROM password enables you to recover a lost boot PROM password by contacting your switch service provider. Without the recovery string, a lost boot PROM password cannot be recovered.

You should set the boot PROM password and the recovery string on all switches, as described in “[With a recovery string](#)” on page 70. If your site procedures dictate that you set the boot PROM password without the recovery string, refer to “[Without a recovery string](#)” on page 72.

## With a recovery string

To set the boot PROM password with a recovery string, refer to the section that applies to your switch model.



**NOTE:** Setting the boot PROM password requires accessing the boot prompt, which stops traffic flow through the switch until the switch is rebooted. You should perform this procedure during a planned downtime.

For the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32, follow this procedure to set the boot PROM password with a recovery string:

1. Connect to the serial port interface as described in “[To connect through the serial port](#)” on page 20.
2. Reboot the switch.
3. Press **ESC** within four seconds after the message **Press escape within 4 seconds** appears.

The following options are available:

- 1 Start system Continues the system boot process
- 2 Recovery password Lets you set the recovery string and the boot PROM password.
- 3 Enter command shell Provides access to boot parameters

4. Enter 2.

If no password was previously set, the following message appears:

Recovery password is NOT set. Please set it now.

If a password was previously set, the following messages appear:

Send the following string to Customer Support for password recovery:

afHTpyLsDo1Pz0Pk5GzhIw==

Enter the supplied recovery password.

Recovery Password:

5. Enter the recovery password (string).

The recovery string must be between 8 and 40 alphanumeric characters. HP recommends a random string that is 15 characters or longer for higher security. The firmware prompts for this password only once. It is not necessary to remember the recovery string because it is displayed the next time you enter the command shell.

The following prompt appears:

New password:

6. Enter the boot PROM password and then reenter it when prompted. The password must be 8 alphanumeric characters (any additional characters are not recorded). Record this password for future use.

The new password is automatically saved (the `saveenv` command is not required).

7. Reboot the switch.

For the Core Switch 2/64 and the SAN Director 2/128, the boot PROM and recovery passwords must be set for each CP card on those switches:

1. Connect to the serial port interface on the standby CP card, as described in “[To connect through the serial port](#)” on page 20.
2. Connect to the active CP card by serial or telnet and enter the `hadisable` command to prevent failover during the remaining steps.
3. For the Core Switch 2/64, reboot the standby CP card by pressing the yellow ejector buttons at top and bottom of the CP card and then pressing both ejector handles back towards the switch to lock the card back into the slot.

For the SAN Director 2/128, reboot the standby CP card by sliding the On/Off switch on the ejector handle of the standby CP card to Off, and then back to On.

4. Press **ESC** within four seconds after the message `Press escape within 4 seconds` appears.

The following options are available:

- 1 Start system                      Continues the system boot process
- 2 Recovery password              Lets you set the recovery string and the boot PROM password
- 3 Enter command shell            Provides access to boot parameters

5. Enter 2.

If no password was previously set, the following message appears:

Recovery password is NOT set. Please set it now.

If a password was previously set, the following messages appear:

Send the following string to Customer Support for password recovery:

`afHTpyLsDo1Pz0Pk5GzhIw==`

Enter the supplied recovery password.

Recovery Password:

6. Enter the recovery password (string).

The recovery string must be between 8 and 40 alphanumeric characters. HP recommends a random string that is 15 characters or longer for higher security. The firmware prompts for this password only once. It is not necessary to remember the recovery string because it is displayed the next time you enter the command shell.

The following prompt appears:

New password:

7. Enter the boot PROM password and then reenter it when prompted. The password must be 8 alphanumeric characters (any additional characters are not recorded). Record this password for future use.

The new password is automatically saved (the `saveenv` command is not required).

8. Connect to the active CP card by serial or telnet and enter the `haenable` command to restore HA, and then fail over the active CP card by entering the `hafailover` command.

Traffic flow through the active CP card resumes when the failover is complete.

9. Connect the serial cable to the serial port on the new standby CP card (previously the active CP card).
10. Repeat step 2 through step 7 for the new standby CP card (each CP card has a separate boot PROM password).
11. Connect to the active CP card by serial or telnet and enter the `haenable` command to restore high availability.

## Without a recovery string

Although you can set the boot PROM password without also setting the recovery string, HP recommends that you set both the password and the string as described in “[With a recovery string](#)” on page 70. If your site procedures dictate that you must set the boot PROM password without the string, follow the procedure that applies to your switch model.



**NOTE:** Setting the boot PROM password requires accessing the boot prompt, which stops traffic flow through the switch until the switch is rebooted. You should perform this procedure during a planned down time.

For the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32, follow this procedure to set the boot PROM password without a recovery string.

1. Create a serial connection to the switch as described in “[To connect through the serial port](#)” on page 20.
2. Reboot the switch by issuing the `reboot` command.
3. Press **ESC** within four seconds after the message `Press escape within 4 seconds` appears.

The following options are available:

- 1 Start system Continues the system boot process
- 2 Recovery password Lets you set the recovery string and the boot PROM password
- 3 Enter command shell Provides access to boot parameters

4. Enter 3.
5. Enter the `passwd` command at the shell prompt.



**NOTE:** The `passwd` command applies to the boot PROM password only when it is entered from the boot interface.

6. Enter the boot PROM password at the prompt and then reenter it when prompted. The password must be 8 alphanumeric characters (any additional characters are not recorded). Record this password for future use.
7. Issue the `saveenv` command to save the new password.
8. Reboot the switch by issuing the `reset` command.

For the Core Switch 2/64 and the SAN Director 2/128, set the password on the standby CP card, fail over, and then set the password on the previously active (now standby) CP card to minimize disruption to the fabric:

1. Determine the active CP card by opening a telnet session to either CP card, connecting as admin, and entering the `hashow` command.
2. Connect to the active CP card by serial or telnet and enter the `hadisable` command to prevent failover during the remaining steps.

3. Create a serial connection to the standby CP card as described in "To connect through the serial port:" on page 20.
4. For the Core Switch 2/64, reboot the standby CP card by pressing the yellow ejector buttons at top and bottom of the CP card and then pressing both ejector handles back towards the switch to lock the card back into the slot.

For the SAN Director 2/128, reboot the standby CP card by sliding the On/Off switch on the ejector handle of the standby CP card to Off, and then back to On.

This causes the card to reset.

5. Press **ESC** within four seconds after the message `Press escape within 4 seconds` appears.

The following options are available:

- 1 Start system              Continues the system boot process
- 2 Recovery password      Lets you set the recovery string and the boot PROM password
- 3 Enter command shell     Provides access to boot parameters

6. Enter 3.

7. Issue the `passwd` command at the shell prompt.



**NOTE:** The `passwd` command only applies to the boot PROM password when it is entered from the boot interface.

8. Enter the boot PROM password at the prompt and then reenter it when prompted. The password must be 8 alphanumeric characters (any additional characters are not recorded). Record this password for future use.
9. Enter the `saveenv` command to save the new password.
10. Reboot the standby CP card by entering the `reset` command.
11. Connect to the active CP card by serial or telnet, enter the `haenable` command to restore HA, and then fail over the active CP card by entering the `hafailover` command.  
Traffic resumes flowing through the newly active CP card after it has completed rebooting.
12. Connect the serial cable to the serial port on the new standby CP card (previously the active CP card).
13. Repeat step 3 through step 10 for the new standby CP card.
14. Connect to the active CP card by serial or telnet and enter the `haenable` command to restore HA.

## Recovering forgotten passwords

If you know the root password, you can use this procedure to recover the user, admin, and factory passwords:

1. Open a CLI connection (serial or telnet) to the switch. If secure mode is enabled, connect to the primary FCS switch.
2. Log in as root.
3. Enter the command for the type of password that was lost:

```
passwd user  
passwd admin  
passwd factory
```

**4.** Enter the requested information at the prompts.

To recover a lost root password, contact your switch service provider.

To recover a lost boot PROM password, contact your switch service provider. You must have previously set a recovery string to recover the boot PROM password.

# 4 Maintaining configurations and firmware

This chapter contains procedures for maintaining switch configurations and installing firmware and consists of the following sections:

- [Maintaining configurations](#), page 75
- [Maintaining firmware](#), page 78
- [Troubleshooting firmware downloads](#), page 86

## Maintaining configurations

It is important to maintain consistent configuration settings on all switches in the same fabric, because inconsistent parameters (such as inconsistent PID formats) can cause fabric segmentation. As part of standard configuration maintenance procedures, HP recommends that you back up all important configuration data for every switch on a host computer server for emergency reference.

The following sections contain procedures for basic switch configuration maintenance.

### Displaying configuration settings

The switch configuration file comprises four sections, and is organized as follows:

- The Boot Parameters section contains variables such as the switch's name and IP address.
- The Licenses section lists the licenses that are active on the switch.
- The Chassis Configuration section contains configuration variables such as diagnostic settings, fabric configuration settings, and SNMP settings.
- The Configuration section contains licensed option configuration parameters.

To display configuration settings, connect to the switch, log in as admin, and enter the `configshow` command at the command line. The configuration settings vary depending on switch model and configuration.

### Backing up a configuration

Keep a backup copy of the configuration file in case the configuration is lost or unintentional changes are made. You should keep individual backup files for all switches in the fabric. You should avoid copying configurations from one switch to another.

The following information is *not* saved in a backup:

- `dnsconfig` information
- passwords

You must have a valid account on the FTP server where the backup file is to be stored.

You can specify the use of secure file copy (scp) during the procedure. For instructions on configuring the use of scp by default, see "[Configuring secure file copy](#)" on page 69.

Before beginning, verify that you can reach the FTP server from the switch. Using a telnet connection, save a backup copy of the configuration file to a host computer as follows:

1. Verify that the FTP service is running on the host computer.
2. Connect to the switch and log in as admin.
3. Enter the `configupload` command.

The command becomes interactive and you are prompted for the required information.

4. Respond to the prompts as follows:

Protocol (scp or ftp)	If your site requires the use of Secure Copy, specify <code>scp</code> . Otherwise, specify <code>ftp</code> .
Server Name or IP Address	Enter the name or IP address of the server where the file is to be stored; for example, <code>192.1.2.3</code> . You can enter a server name if DNS is enabled.
User name	Enter the user name of your account on the server; for example, <code>JohnDoe</code> .
File name	Specify a file name for the backup file; for example, <code>config.txt</code> . Use the forward slash (/) to specify absolute path names. Relative path names create the file in the user's home directory on UNIX servers, and in the directory where the FTP server is running on Windows servers.
Password	Enter your account password for the server.

Example:

```
switch:admin> configupload
Protocol (scp or ftp) [ftp]: ftp
Server Name or IP Address [host]: 192.1.2.3
User Name [user]: JohnDoe
File Name [config.txt]: /pub/configurations/config.txt
Password: *****
Upload complete
switch:admin>
```

## Restoring a configuration

Restoring a configuration involves overwriting the configuration on the switch by downloading a previously saved backup configuration file. Perform this procedure during a planned down time.

Make sure that the configuration file you are downloading is compatible with your switch model, because configuration files from other model switches might cause your switch to fail.

You must have a user ID on the FTP server where the backup file is stored.

Use the following procedure:

1. Verify that the FTP service is running on the server where the backup configuration file is located.
2. Connect to the switch and log in as admin.
3. Disable the switch by entering the `switchdisable` command.
4. Enter the `configdownload` command.

The command becomes interactive and you are prompted for the required information.

**5.** Respond to the prompts as follows:

Protocol (scp or ftp)	If your site requires the use of Secure Copy, specify <code>scp</code> . Otherwise, specify <code>ftp</code> .
Server Name or IP Address	Enter the name or IP address of the server where the file is stored; for example, <code>192.1.2.3</code> . You can enter a server name if DNS is enabled.
User name	Enter the user name of your account on the server; for example, <code>JohnDoe</code> .
File name	Specify the full path name of the backup file; for example, <code>/pub/configurations/config.txt</code> .
Password	Enter your account password for the server.

**6.** At the `Do you want to continue [y/n]` prompt, enter `y`.

**7.** Wait for the configuration to be restored.

**8.** When the process is finished, enter the `switchenable` command:

Example:

```
switch:admin> configdownload
Protocol (scp or ftp) [ftp]: ftp
Server Name or IP Address [host]: 192.1.2.3
User Name [user]: JohnDoe
File Name [config.txt]: /pub/configurations/config.txt
Password: xxxxxx
*** CAUTION ***
This command is used to download a backed-up configuration
for a specific switch. If using a file from a different
switch, this file's configuration settings will override
any current switch settings. Downloading a configuration
file, which was uploaded from a different type of switch,
may cause this switch to fail.

Do you want to continue [y/n]: y
download complete..
switch:admin> switchenable
```



**NOTE:** Because some configuration parameters require a reboot to take effect, after you download a configuration file you must reboot to be sure that the parameters are enabled. Before the reboot, this type of parameter is listed in the configuration file, but it is not effective until after the reboot.

## Downloading configurations across a fabric

To save time when configuring fabric parameters and software features, you can save a configuration file from one switch and download it to other switches of the same model type, as described in the following procedure. Avoid downloading configuration files to different model switches, because that can cause the switches to fail.

1. Configure one switch first.
2. Use the `configupload` command to save the configuration information. See “[Backing up a configuration](#)” on page 75.
3. Use the `configdownload` command to download the configuration to each of the remaining switches. See “[Restoring a configuration](#)” on page 76.

## Editing configuration files

Beginning with Fabric OS v4.2.0, the `portcfg` line in the configuration file for a brand new switch contains 256 entries, regardless of the number of ports on the switch. This line length exceeds the capacity of the vi editor. If you must edit a new configuration file, you can do so with the vim editor. Or, be sure to perform a `portcfg` operation before attempting to edit the configuration file (because after the `portcfg` operation, the `portcfg` line in the configuration file contains only as many entries as the maximum number of ports on the switch).

## Printing hard copies of switch information

HP recommends that you print a hard copy of all key configuration data, including license key information for every switch, and store it in a safe and secure place for emergency reference. Print out the information from the following commands, and store the printouts in a secure location:

- `configshow` displays configuration parameters and setup information, including license information.
- `ipaddrshow` displays the IP address.
- `licenseshow` displays the license keys you have installed and provides better detail than the license information from the `configshow` command.

Depending on the security procedures of your company, you might also want to keep a record of the user levels and passwords for all switches in the fabric. Access to this sensitive information should be limited.

---

## Maintaining firmware

This section explains how to obtain and install firmware. Fabric OS v4.4.0 provides nondisruptive firmware installation.

In most cases, you are to upgrade firmware; that is, install a newer firmware version than the one you are currently running. However, some circumstances may require installing an older version; that is, *downgrading* the firmware. The procedures in this section assume that you are upgrading firmware, but they work for downgrading as well, provided the old and new firmware versions are compatible.

Using the CLI (or HP Advanced Web Tools), you can upgrade the firmware on one switch at a time. You can use the optionally licensed HP Fabric Manager software tool to upgrade firmware simultaneously on multiple switches. For more details on Fabric Manager and other licensed software tools, go to the HP StorageWorks web site: <http://www.hp.com/country/us/eng/prodserv/storage.html>.

## Obtaining and unzipping firmware

Firmware upgrades are available for customers with support service contracts and partners on the HP StorageWorks web site: <http://www.hp.com/country/us/eng/prodserv/storage.html>.

The firmware is delivered in a compressed file that contains RPM packages with names defined in a `pfile`, a binary file that contains specific firmware information (timestamp, platform code, version, and so forth) and the names of the packages of firmware to be downloaded. You must unzip the firmware (using the UNIX `tar` or `gzip` command, or a Windows `unzip` program) before you can use the `firmwaredownload` command to update the firmware on your equipment.

When you unpack the downloaded firmware it expands into a directory that is named according to the version of Fabric OS it contains. For example, if you download and unpack Fabric OS v4.4.0.zip, it expands into a directory called `v4.4.0`. When you use the `firmwaredownload` command, you specify the path to the `v4.4.0` directory and append the keyword `release.plist` to the path.

## Checking connected switches

If the switch to be upgraded is running v4.1.0 firmware (or later), HP recommends that all switches directly connected to it be running versions no earlier than v2.6.1, v3.1.0, or v4.1.0. If some connected switches are running older firmware, upgrade them to at least the earliest recommended version (shown in [Table 11](#)) before upgrading firmware on your switch.

**Table 11** Recommended firmware

HP StorageWorks switch <sup>1</sup>	Earliest recommended Fabric OS version
1 GB	v2.6.1
SAN Switch 2/8-EL, SAN Switch 2/16-EL, SAN Switch 2/16	v3.1.0
SAN Switch 2/8V, SAN Switch 2/16V	v4.2.0
SAN Switch 2/32	v4.1.0
SAN Switch 4/32	v4.4.0
Core Switch 2/64	v4.1.0
SAN Director 2/128	v4.2.0

1. During code activation on 2 GB switches, SAN Switch 2/8V, SAN Switch 2/16V, or SAN Switch 2/32 running Fabric OS v4.1.0 or later, data continues to flow between hosts and storage devices; however, fabric services are unavailable for a period of approximately 50-55 seconds. Possible disruption of the fabric can be minimized by ensuring that switches logically adjacent to these models (directly connected via an ISL) are running at the minimum Fabric OS v2.6.1 or later, v3.1.0 or later, or v4.1.0 or later. If 2 GB switches, SAN Switch 2/8V, SAN Switch 2/16V, or SAN Switch 2/32 are adjacent and you start firmware downloads on them at the same time, I/O might be disrupted.

To determine whether you need to upgrade connected switches before upgrading your switch, use the following procedure on each connected switch to display firmware information and build dates.

1. Connect to the switch and log in as admin.
2. Issue the `version` command.

The following information is displayed:

- Kernel: Displays the version of switch kernel operating system
- Fabric OS: Displays the version of switch Fabric OS
- Made on: Displays the build date of firmware running in switch
- Flash: Displays the install date of firmware stored in nonvolatile memory
- BootProm: Displays the version of the firmware stored in the boot PROM

## About the download process

The `firmwareDownload` command downloads unzipped switch firmware from an FTP server to the switch's nonvolatile storage area.

In the Core Switch 2/64 and SAN Director 2/128, this command by default downloads the firmware image to the two CP cards in rollover mode, to prevent disruption to application services. This operation depends on HAHA support. If HA is not available, experienced technicians can upgrade the CPs one at a time, using the `-s` option.

HP StorageWorks fixed-port switches and each CP card of the Core Switch 2/64 and SAN Director 2/128 have two partitions of nonvolatile storage areas (a primary and a secondary) to store two firmware images. The `firmwaredownload` command always loads the new image into the secondary partition and swaps the secondary partition to be the primary. It then reboots the partition and activates the new image. Finally, it performs the `firmwarecommit` procedure automatically, to copy the new image to the other partition.

## Effects of firmware changes on accounts and passwords

**Table 12** describes what happens to accounts and passwords when you replace the switch firmware with a different version. *Upgrading* means installing a newer version of firmware. *Downgrading* means installing an older version of firmware.

**Table 12** Effect of firmware on accounts and passwords

Change	First time	Subsequent times (after upgrade, then downgrade, then upgrade)
Upgrading	Default accounts and their passwords are preserved.	User-defined and default accounts and their passwords are preserved.
Downgrading	User-defined accounts are no longer valid. Default accounts and their passwords are preserved. If a default account was disabled, it is reenabled after the downgrade.	User-defined and default accounts and their passwords are preserved, including accounts added after the first upgrade.
Upgrading to v3.2.0	(You may upgrade a switch in the fabric as part of "Checking connected switches" on page 79.) Earlier versions allowed you to change the default account names. You cannot add user-defined accounts until you change the names back to default with the <code>passwdDefault</code> command.	

For more details on older releases of Fabric OS, see "Understanding legacy password behavior" on page 229.

## Considerations for downgrading firmware

The following items must be considered before attempting to downgrade to an earlier version of Fabric OS:

- If your fabric is set to the extended edge PID format and you want to downgrade to an older Fabric OS version that does not support extended edge, you must change the PID to a supported format. For more information, see "Configuring the PID format" on page 203.
- Downgrading a SAN Director 2/128 that is configured for two domains from Fabric OS v4.4.0 to Fabric OS v4.2.0 is not supported.
- If you are running v4.0.2 firmware on a SAN Switch 2/32, you cannot downgrade to earlier versions.

## Upgrading HP StorageWorks SAN switches

SAN StorageWorks Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32 maintain primary and secondary partitions for firmware. The `firmwaredownload` command defaults to an Auto Commit option that automatically copies the firmware from one partition to the other.

Do not override Auto Commit under normal circumstances; use the default. If you override the Auto Commit option (that is, use the single mode `-s` option with the `firmwaredownload` command and then specify `no` to the Auto Commit prompt), and then reboot with the `hareboot` command, you must execute the `firmwarecommit` command.

As an option, before starting a firmware download, HP suggests that you connect the switch with a console cable to a computer that is running a session capture. The information collected may be useful if needed for troubleshooting.

## Summary of the upgrade process

The following summary describes the default behavior of the `firmwareDownload` command (without options) on the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32.

- Issue the `firmwareDownload` command.

Fabric OS downloads firmware to the secondary partition.

The system performs an HA reboot (`hareboot`). After the `hareboot`, the former secondary partition is now the primary partition. The system replicates the firmware from the primary to the secondary partition.

- Issue the `firmwareDownloadStatus` command to view the firmware process.

## SAN Switch upgrade procedure

For the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32, the upgrade process first downloads and then commits the firmware to the switch. While the upgrade is proceeding, you can start another telnet session on the switch and observe the upgrade progress if you wish.



**NOTE:** After you start the process, do not enter any disruptive commands (such as `reboot`) that interrupt the process. The firmware download and commit process takes approximately 15 minutes. If there is a problem, wait for the time-out (30 minutes for network problems; 10 minutes for incorrect IP address). Disrupting the process can render the switch inoperable and require you to seek help from Customer Support.

Do not disconnect the switch from power during the process; the switch could become inoperable upon reboot.

---

Use the following procedure to upgrade firmware for the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32:

1. Verify that the FTP service is running on the host server and that you have a user ID on that server.
2. Obtain the firmware file from the HP StorageWorks web site at <http://www.hp.com/country/us/eng/prodserv/storage.html> and store the file on the FTP server. Verify that the FTP service is running.
3. Issue the `firmwareShow` command to check the current firmware version on connected switches. Upgrade their firmware if necessary before proceeding with upgrading this switch. See "Checking connected switches" on page 79.
4. Connect to the switch and log in as admin.
5. Issue the `firmwareShow` command to check the current firmware version of the switch to verify compatibility with the version of firmware you are going to download.



**NOTE:** For the SAN Switch 2/32, if you are running Fabric OS v4.0.2, you cannot downgrade to earlier versions.

For the SAN Switches 2/8V and 2/16V, if you are running Fabric OS v4.2.0, you cannot downgrade to earlier version.

---

6. Issue the `firmwareDownload` command.

**7.** At the Do you want to continue [y/n] prompt, enter y.

**8.** Respond to the prompts as follows:

Server Name Enter the name or IP address of the server where the firmware file is stored; or IP Address: for example, 192.1.2.3. You can enter a server name if DNS is enabled.

User name: Enter the user name of your account on the server; for example, JohnDoe.

File name: Specify the full path name of the firmware directory, appended by release.plist; for example, /pub/v4.4.0/release.plist.

Password: Enter your account password for the server.

After the firmware is downloaded, the switch reboots and starts the firmware commit.

**9.** After the reboot, connect to the switch and log in again as admin.

**10.** If you want to watch the upgrade progress, issue the firmwaredownloadstatus command to monitor the status of the firmware download.

**11.** After the firmware commit finishes, issue the firmwareshow command to display the firmware level for both partitions.

Example:

```
switch:admin> firmwaredownload
You can run firmwareDownloadStatus to get the status of this
command.
This command will cause the switch to reset and will require that
existing telnet, secure telnet or SSH sessions be restarted.
Do you want to continue [Y]: y
Server Name or IP Address: 192.1.2.3
User Name: JohnDoe
File Name: /pub/v4.4.0/release.plist
Password: xxxxxx
Firmwaredownload has started.

0x8fd (Fabric OS): Switch: 0, Warning SULIB-FWDL_START, 3,
Firmwaredownload command has started.
.
.
.
```

Log in again to view the upgrade progress; for example:

```
switch:admin> firmwaredownloadstatus
[0]: Tue Apr 20 10:32:34 2004
cp0: Firmwaredownload has started.
[1]: Tue Apr 20 10:36:07 2004
cp0: Firmwaredownload has completed successfully.
[2]: Tue Apr 20 10:57:09 2004
cp0: Firmwarecommit has started.
[3]: Tue Apr 20 10:36:07 2004
cp0: Firmwarecommit has completed successfully.
[4]: Tue Apr 20 11:03:28 2004
cp0: Firmwaredownload command has completed successfully.
switch:admin> firmwareshow
Primary partition: v4.4.0
Secondary Partition: v4.4.0
switch:admin>
```

# Upgrading the Core Switch 2/64 and the SAN Director 2/128

You can download firmware to the Core Switch 2/64 and SAN Director 2/128 without disrupting the overall fabric if the two CP cards are installed and fully synchronized. Use the `hashow` command to confirm synchronization. If only one CP card is powered on, the switch must reboot to activate firmware, which is disruptive to the overall fabric.

If there is an error during the firmware download, the system ensures that the two partitions of a CP card contain the same version of firmware. However, the two CP cards might contain different versions of firmware; in that event, repeat the firmware download process.

During the upgrade process the director fails over to its standby CP card and the IP addresses for the two logical switches move to that CP card's Ethernet port. This might cause informational ARP address reassignment messages to appear on other switches in the fabric. This is normal behavior, because the association between the IP addresses and MAC addresses has changed.

## Summary of the upgrade process

The following summary describes the default behavior of the `firmwaredownload` command (without options) on the Core Switch 2/64 and SAN Director 2/128.

- Issue the `firmwaredownload` command on the active CP card.
- The standby CP card downloads firmware.
- The standby CP card reboots and comes up with the new Fabric OS.
- The active CP card synchronizes its state with the standby CP card.
- The active CP card forces a failover and reboots to become the standby CP card.
- The *new* standby CP card (the active CP card before the failover) downloads firmware.
- The *new* standby CP card reboots and comes up with the new Fabric OS.
- The *new* active CP card synchronizes its state with the *new* standby CP card.
- The `firmwarecommit` command runs automatically on both CP cards.



**NOTE:** After you start the process, do not enter any disruptive commands (such as `reboot`) that interrupt the process. The entire firmware download and commit process takes approximately 15 minutes. If there is a problem, wait for the time-out (30 minutes for network problems; 10 minutes for incorrect IP address). Disrupting the process can render the switch inoperable and require you to seek help from Customer Support.

Do not disconnect the switch from power during the process, because the switch could become inoperable upon reboot.

## Core Switch 2/64 and SAN Director 2/128 upgrade procedure

The Core Switch 2/64 has four IP addresses: one for each of the two logical switches (switch 0 and switch 1) and one for each of the two CP cards (CP0 in slot 5 and CP1 in slot 6). The SAN Director 2/128 in its default configuration has three IP addresses, but it can be configured for four.



**NOTE:** By default, the `firmwaredownload` command automatically upgrades both the active CP card and the standby CP card. When upgrading a Core Switch 2/64 that is running v4.0.0c or earlier, you must upgrade each CP card separately, as described in "To upgrade a single Core Switch 2/64 or SAN Director 2/128 CP card:" on page 240. You should not use this procedure under normal circumstances.

Follow this procedure to upgrade the firmware on the Core Switch 2/64 and the SAN Director 2/128:

1. Verify that the FTP service is running on the host server and that you have a user ID on that server.
2. Obtain the firmware file from the HP StorageWorks web site at <http://www.hp.com/country/us/eng/prodserv/storage.html> and store the file on the FTP server. Verify that the FTP service is running.
3. Use the `firmwareshow` command to check the current firmware version on connected switches. Upgrade their firmware if necessary before proceeding with upgrading this switch.  
See "Checking connected switches" on page 79.
4. Using a telnet session, connect to the switch and log in as admin.
5. For the Core Switch 2/64, issue the `firmwareshow` command to check the current firmware version of the switch.

If the switch is running v4.0.0c or earlier, and you want to downgrade to an earlier version, you must load firmware to each CP card separately using the procedure in "To upgrade a single Core Switch 2/64 or SAN Director 2/128 CP card:" on page 240.

6. Issue the `hashow` command to confirm that the two CP cards are synchronized. CP cards must be synchronized and running Fabric OS v4.1.0 or later to provide a nondisruptive download. If the two CP cards are not synchronized, and the current firmware version is 4.1.0 or later, issue the `hasyncstart` command to synchronize the two CP cards. In the following example, the active CP card is CP1 and the standby CP card is CP0.

Example:

```
switch:admin> hashow
Local CP (Slot 6, CP1): Active
Remote CP (Slot 5, CP0): Standby
HA Enabled, Heartbeat up, HA State is in Sync
switch:admin>
```

7. Log in to either of the logical switches.
8. Issue the `firmwaredownload` command.
9. At the Do you want to continue [y/n] prompt enter: y

**10.** Respond to the prompts as follows:

Server Name	Enter the name or IP address of the server where the firmware file is stored; for example, 192.1.2.3. You can enter a server name if DNS is enabled.
User name	Enter the user name of your account on the server; for example, JohnDoe.
File name	Specify the full path name of the firmware directory, appended by release.plist; for example, /pub/v4.4.0/release.plist.
Password	Enter your account password for the server.

The firmware is downloaded to one CP card at a time, beginning with the standby CP card. During the process, the active CP card is failed over. After the firmware is downloaded, a firmware commit starts on both CP cards.

**11.** Optionally, after the failover, connect to the switch and log in again as admin.

**12.** Issue the firmwaredownloadstatus command to monitor the firmwaredownload status.

**13.** Issue the firmwareshow command to display the new firmware versions.

Example:

```
switch:admin> firmwaredownload
This command will upgrade both CPs in the switch. If you
what to upgrade a single CP only, please use -s option.

You can run firmwareDownloadStatus to get the status
of this command.

This command will cause the active CP to reset and will
require that existing telnet, secure telnet, or SSH sessions
be restarted.

Do you want to continue [Y]: y
Server Name or IP Address: 192.1.2.3
User Name: JohnDoe
File Name: /pub/v4.4.0/release.plist
Password:*****
FirmwareDownload has started on Standby CP. It may take up to 30
minutes.
Firmwaredownload has completed successfully on Standby CP.

.
.
.

Standby CP reboots.
Standby CP booted up.
Standby CP booted up with new firmware.
cp1: Firmwarecommit has started on both Active and Standby CPs.
cp1: Firmwarecommit has completed successfully on Active CP.
cp1: Firmwaredownload command has completed successfully.

switch:admin>
```

Start a new session to view the upgrade progress:

```
switch:admin> firmwaredownloadstatus
[0]: Tue Apr 20 15:18:56 2003
cp0: Firmwaredownload has started on Standby CP. It may take up to 10
minutes.
[1]: Tue Apr 20 15:24:17 2003
cp0: Firmwaredownload has completed successfully on Standby CP.
[2]: Tue Apr 20 15:24:19 2003
cp0: Standby CP reboots.
[3]: Tue Apr 20 15:27:06 2003
cp0: Standby CP booted up.
[4]: Tue Apr 20 15:29:01 2003
cp1: Active CP forced failover succeeded. Now this CP becomes Active.
[5]: Tue Apr 20 15:29:05 2003
cp1: Firmwaredownload has started on Standby CP. It may take up to 30
minutes.
[6]: Tue Apr 20 15:34:16 2003
cp1: Firmwaredownload has completed successfully on Standby CP.
[7]: Tue Apr 20 15:34:19 2003
cp1: Standby CP reboots.
[8]: Tue Apr 20 15:36:59 2003
cp1: Standby CP booted up with new firmware.
[9]: Tue Apr 20 15:37:04 2003
cp1: Firmwarecommit has started on both Active and Standby CPs.
[10]: Tue Apr 20 15:42:48 2003
cp1: Firmwarecommit has completed successfully on Active CP.
[11]: Tue Apr 20 15:42:49 2003
cp1: Firmwaredownload command has completed successfully.
```

## Troubleshooting firmware downloads

A firmware download can fail for many reasons, such as a power failure, a failed network connection, a failed FTP server, or an incorrect path to unpacked firmware files. In most cases, the firmware is not affected. You can make necessary corrections (for example, check the Ethernet cables and check the file path names) and then run the `firmwaredownload` command again.



**NOTE:** Under firmware versions earlier than v4.1.0, do not perform a firmware download while the switch is running POST. If a firmware download is attempted on a Core Switch 2/64 while POST is running, it might fail because the CP cards cannot synchronize with each other.

Issue the `firmwareshow` command to see whether both CP cards have the same firmware. In this example, the active CP card has the old version of firmware and the standby CP card has the new version:

```
switch: admin> firmwareshow
Local CP (Slot 5, CP0): Active
    Primary partition: v4.2.0
    Secondary Partition: v4.2.0
Remote CP (Slot 6, CP1): Standby
    Primary partition: v4.4.0
    Secondary Partition: v4.4.0
switch: admin>
```

Decide which firmware version you want to be applied to both CP cards. Then repeat the download procedure.



# 5 Configuring the Core Switch 2/64 and the SAN Director 2/128

This chapter contains procedures that are specific to the Core Switch 2/64 and the SAN Director 2/128 and consists of the following sections:

- [Identifying ports](#), page 89
- [Basic card management](#), page 90
- [Setting chassis configurations](#), page 92
- [Setting the card beacon mode](#), page 98

Because these switches contain interchangeable 16-port cards (the software calls them *blades*), their procedures differ from those for the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and the SAN Switch 4/32 fixed-port switches. For example, fixed-port models identify ports by *domain, port number*, while director models identify ports by *slot/port number*.

Also, because the Core Switch 2/64 comprises two logical switches (*domains*) and the SAN Director 2/128 in its default configuration has only one domain, procedures for the two directors sometimes differ from one another.

For detailed information about the Core Switch 2/64 and the SAN Director 2/128, refer to their hardware reference manuals.

## Identifying ports

The Core Switch 2/64 and the SAN Director 2/128 have slots and can have a variable number of ports within a given domain. Ports are identified by their combined slot number and port number.

There are a total of 10 slots that contain cards:

- Slot numbers 5 and 6 contain *control processor* cards (CPs).
- Slot numbers 1 through 4 and 7 through 10 contain port cards.

On each port card, there are 16 ports (counted from the bottom, 0 to 15). A particular port must be represented by both slot number (1 through 4 and 7 through 10) and port number (0 through 15).

The Core Switch 2/64 is divided into two logical switches, where slots 1 through 4 are logical switch 0 (sw0) and slots 7 through 10 are logical switch 1 (sw1). You must be connected to the logical switch that represents the slot where you want to execute a command.

In the SAN Director 2/128 default configuration, all the ports are part of a single logical switch. With Fabric OS v4.4.0 and later, you can configure the SAN Director 2/128 as two logical switches (*domains*).

The following sections tell how to identify ports on the Core Switch 2/64 and the SAN Director 2/128, and how to identify ports for zoning commands.

## By slot and port number

To select a specific port in the Core Switch 2/64 and the SAN Director 2/128, you must identify both the slot number and the port number using the format *slot number/port number*. No spaces are allowed between the slot number, the slash (/), and the port number.

The following example shows how to enable port 4 on a card in slot 2:

```
switch:admin> portenable 2/4
```

## By port area ID

Zoning commands require that you specify ports using the area ID method. In Fabric OS v4.0.0 and later, each port on a particular domain is given a unique area ID. How the port number is related to the area ID depends upon the PID format used in the fabric:

- When Core PID mode is in effect, the area ID for port 0 is 0, for port 1, it is 1, and so forth. When using Core PID mode on the Core Switch 2/64 (two logical 64-port switches) and the SAN Director 2/128 configured with two domains, the area IDs for both logical switches (domains) range from 0 to 63. This means that both logical switch 0 and logical switch 1 have a port that is referenced with area ID 0.
- When Extended Edge PID mode is in effect, the area ID is the port number plus 16 for ports 0 to 111. For port numbers higher than 111, the area ID wraps around so that port 112 has an area ID of 0, and so on. Each 64-port logical switch (domain) has area IDs ranging from 16 to 79.

To determine the area ID of a particular port, enter the `switchshow` command. This command displays all ports on the current (logical) switch and their corresponding area IDs.

---

## Basic card management

The following sections provide procedures for powering a card on and off and for disabling and enabling a card.

### Powering port cards on and off

Port cards are powered on by default.

To power off a port card:

1. Connect to the switch and log in as admin.
2. Issue the `slotpoweroff` command with the slot number of the card you want to power off.

The slot must exist in the logical switch where you are logged in.

Example:

```
switch:admin> slotpoweroff 3
Slot 3 is being powered off
switch:admin>
```

To provide power to a port card:

1. Connect to the switch and log in as admin.
2. Issue the `slotpoweron` command with the slot number of the card you want to power on.

The slot must exist in the logical switch where you are logged in.

Example:

```
switch:admin> slotpoweron 3
Powering on slot 3
switch:admin>
```

## Disabling and enabling cards

Cards are enabled by default.

You might need to disable a card to perform diagnostics. When diagnostics are executed manually (from the Fabric OS command line), many commands require the card to be disabled. This ensures that diagnostic activity does not interfere or disturb normal fabric traffic.

To disable a card:

1. Connect to the switch and log in as admin.
2. Issue the `slotoff` command with the slot number of the card you want to disable.

Example:

```
switch:admin> slotoff 3
Slot 3 is being disabled
switch:admin>
```

To enable a card:

1. Connect to the switch and log in as admin.
2. Issue the `sloton` command with the slot number of the card you want to enable.

Example:

```
switch:admin> sloton 3
Slot 3 is being enabled
switch:admin>
```

## Conserving power

To conserve power and ensure that more critical components are the least affected by a power fluctuation, you can power off components in a specified order using the `powerofflistset` command.

The available power is compared to the power demand to determine if there is enough power to operate. If there is less power available than the demand, the power-off list is processed until there is enough power for operation. By default, the processing proceeds from slot 1 to the last slot in the chassis. As power becomes available, slots are powered up in the reverse order.



**NOTE:** Some FRUs in the chassis may use significant power, yet cannot be powered off through software. For example, a missing blower FRU may change the power computation enough to affect how many slots can be powered up.

The `powerofflistshow` command displays the power-off order.

# Setting chassis configurations

The `chassisconfig` command allows you to set the chassis configuration for products that support both single-switch (one domain) and dual-switch (two domains) operation.

**Table 13** lists the supported options for Fabric OS v4.4.0 or later. In the table, Blade ID 4 indicates a SAN Director 2/128 card, and Blade ID 2 indicates a Core Switch 2/64 card.

**Table 13** Supported options

Option	Result
1	One 128-port switch (Blade ID 4 on slots 1–4, 7–10)
2	Two 64-port switches (Blade ID 4 on slots 1–4, 7–10)
3	Two 64-port switches (Blade ID 4 on slots 1–4, ID 2 on slots 7–10)
4	Two 64-port switches (Blade ID 2 on slots 1–4, ID 4 on slots 7–10)

The following sections contain procedures for obtaining chassis information, and for configuring director domains using the `chassisconfig` command.

## Obtaining slot information

For a Core Switch 2/64 or a SAN Director 2/128 configured as two logical switches, the chassis-wide commands display or control both logical switches. In the default configuration, the SAN Director 2/128 is configured as one logical switch, so the chassis-wide commands display and control the single logical switch.

To display the status of all slots in the chassis:

1. Connect to the switch and log in as user or admin.
2. Issue the `slotshow` command to display the current status of each slot in the system.

The format of the display includes a header and four fields for each slot. The fields and their possible values are:

**Table 14** Header fields

Field	Value
Slot	Displays the physical slot number.
Blade type	Displays the card type: <ul style="list-style-type: none"><li>• SW BLADE: The card is a switch.</li><li>• CP BLADE: The card is a control processor.</li><li>• UNKNOWN: The card is not present or its type is not recognized.</li></ul>

**Table 14** Header fields (continued)

Field	Value
ID	Displays the hardware ID of the card type
Status	Displays the status of the card: <ul style="list-style-type: none"><li>• VACANT: the slot is empty.</li><li>• INSERTED, NOT POWERED ON: The card is present in the slot, but is turned off.</li><li>• DIAG RUNNING POST1: The card is present, powered on, and running the post initialization power on self tests.</li><li>• DIAG RUNNING POST2: The card is present, powered on, and running the power-on self test.</li><li>• ENABLED: The card is on and enabled.</li><li>• ENABLED (User Ports Disabled): The card is on, but external ports have been disabled with the <code>bladedisable</code> command.</li><li>• DISABLED: The card is powered on, but disabled.</li><li>• FAULTY: The card is faulty because an error has been detected. The reason code numbers displayed are for use in debugging.</li><li>• UNKNOWN: The card is inserted but its state cannot be determined.</li></ul>

## Configuring a new SAN Director 2/128 with two domains

By default, the SAN Director 2/128 is configured as one 128-port switch (one domain). Use the following procedure to add a new SAN Director 2/128 to a fabric and configure it as two 64-port switches (two domains). The procedure assumes that the new director:

- Has been installed and connected to power, but is not yet attached to the fabric.
- Has been given an IP address, but is otherwise running factory defaults.

If this is not the case, back up the current configuration before starting, so that you can restore it later if necessary.

- Is running Fabric OS v4.4.0 or later.

1. Connect to the switch and log in as admin.

2. Issue the `chassisconfig` command without options to verify that the switch is configured with one domain. For example:

```
chassisconfig
```

```
Current Option: 1
```

3. Issue the `chassisconfig` command to configure two domains. Use the `-f` option to suppress prompting for uploading the configuration. This command reboots the system.

```
chassisconfig -f 2
```

```
Current Option changed to 2
```

```
Restoring switch 0 configuration to factory defaults...
```

```
All account passwords have been successfully set to factory default.
```

```
Restoring switch 1 configuration to factory defaults...
```

```
All account passwords have been successfully set to factory default.
```

4. After the system reboots, log in again to the first logical switch (sw0) as admin.
5. Use the `configure` command to configure the sw0 to match your fabric specifications.

If the director is to be merged into an existing fabric, do not configure zoning parameters; these are propagated automatically when you merge the director into the fabric.
6. Log in to the second logical switch (sw1) as admin.
7. Use the `configure` command to configure the sw1 to match your fabric specifications.

If the director is to be merged into an existing fabric, do not configure zoning parameters; these are propagated automatically when you merge the director into the fabric.
8. If the fabric is in secure mode, perform the following steps; otherwise, proceed to [step 9](#). (Refer to the *HP StorageWorks Secure Fabric OS user guide* for specific instructions.)
  - a. Optionally, to configure sw0 and sw1 in one operation, connect them with an ISL link to form a temporary fabric.
  - b. If you want sw0 and sw1 to be fabric configuration servers, update the overall fabric's FCS policy to include them. If not, skip this step.
  - c. On sw0, enable security mode and use the `secmodeenable` command to create an FCS list that matches your overall fabric's FCS policy.
  - d. Reset the version stamp on sw0.
  - e. If you connected sw0 and sw1 in [step 8a](#) and you do not want them connected, disconnect the ISL link between them. If you did not connect them, repeat [step 8b](#) through [step 8d](#) on sw1.
9. Optional: Connect the new two-domain SAN Director 2/128 to the fabric.
10. Issue the `fabricshow` command to verify that sw0 and sw1 have been merged with the fabric.
11. Issue the `cfgshow` command to verify that zoning parameters were propagated.

## Converting an installed SAN Director 2/128 to support two domains

Fabric OS versions earlier than v4.4.0 supported only one domain for the SAN Director 2/128 (one 128-port logical switch). When you upgrade a SAN Director 2/128 to Fabric OS v4.4.0 or later, you can use the `chassisconfig` command to specify two domains for the director (two 64-port logical switches, sw0 and sw1).



**NOTE:** This procedure restores most configuration parameters to factory defaults. After performing this procedure, you must check the new configuration and reconfigure those parameters that you customized in the old configuration.

During this procedure, power is reset and the CP cards are rebooted, so traffic on the fabric is disrupted. If the fabric is in secure mode, enabling security on the new domains is a complicated task. You should avoid converting existing core switches.

1. Connect to the switch and log in as admin.
2. If the director is already in a fabric, minimize disruption by removing the director from the fabric using one of the following methods:
  - Physically disconnect the director.
  - Use the `portcfgpersistentdisable` command on all connected remote switches to persistently disable ports that are connected to the director.

3. Issue the `chassisconfig` command to change the configuration from the default (one domain) to two domains. This command reboots the system.

```
chassisconfig 2
```

During the conversion, you are prompted to save the configuration of sw0. Follow the prompts to save the configuration file.

4. After the system reboots, log in again as admin to each logical switch.
5. Using the configuration file saved in [step 3](#) as a guide, manually reconfigure sw0 and sw1.

Do not configure zoning parameters; these are propagated automatically when you merge the director into the fabric.
6. If the fabric is in secure mode, perform the following steps; otherwise, proceed to [step 7](#).
  - a. Optionally, to configure sw0 and sw1 in one operation, connect them with an ISL link to form a temporary fabric.
  - b. If you want sw0 and sw1 to be fabric configuration servers, update the overall fabric's FCS policy to include them. If not, skip this step.
  - c. On sw0, enable security mode and use the `secmodeenable` command to create an FCS list that matches your overall fabric's FCS policy.
  - d. Reset the version stamp on sw0.
  - e. If you connected sw0 and sw1 in [step 6a](#) and you do not want them connected, disconnect the ISL link between them. If you did not connect them, repeat [step 6b](#) through [step 6d](#) on sw1.
7. If you physically disconnected the switch in [step 2](#), reconnect it to the fabric.

If you used the `portcfgpersistentdisable` command in [step 2](#), use the `portcfgpersistenteable` command to persistently enable all ports that connect the switch to other switches in the fabric.
8. Use the `fabricshow` command to verify that sw0 and sw1 have been merged with the fabric.
9. Use the `configshow` command to verify that zoning parameters were propagated.

## Combining Core Switch 2/64 and SAN Director 2/128 cards in one chassis

You can preserve your investment in legacy equipment by combining Core Switch 2/64 cards and SAN Director 2/128 cards in one chassis.

The following procedure assumes that:

- The Core Switch 2/64 has one logical switch (sw0, slots 1 through 4) populated with port cards. (You can perform the same procedure on sw1 slots 7–10.) The other side of the chassis is empty.
- Fabric OS firmware v4.4.0 or later is already installed on the new SAN Director 2/128 CP cards.

The result of the procedure is a system populated with four Core Switch 2/64 port cards in slots 1 through 4, two SAN Director 2/128 CP cards in slots 5 and 6, and four SAN Director 2/128 port cards in slots 7 through 10 and configured with two domains.

Consider the following rules and guidelines:

- Because this procedure requires power reset and rebooting, traffic on the fabric is disrupted.
- You should be familiar with the standard procedures for shutting down the equipment. Refer to the *HP StorageWorks Core Switch 2/64 and SAN Director 2/128 installation guide*, which contains more details on disconnecting an HP StorageWorks model from the network and fabric.

- The result of this procedure is two 64-port logical switches (domains) that communicate through external ISLs.
- Only similar port cards can be inserted in the same logical switch (slots 1 through 4 or slots 7 through 10); you cannot install Core Switch 2/64 and SAN Director 2/128 port cards in the same logical switch.
- Before installing Core Switch 2/64 cards in a SAN Director 2/128 chassis, review the power supply requirements in the Core Switch 2/64 hardware reference manual and make sure you meet the higher power requirements of the Core Switch 2/64 cards. You need enough power supplies in the SAN Director 2/128 chassis to ensure uninterrupted performance if a power supply fails.
- You must replace both of the Core Switch 2/64 CP cards with SAN Director 2/128 CP cards running Fabric OS v4.4.0 or later. Using dissimilar CP cards in the same chassis is not allowed.

To combine Core Switch 2/64 and SAN Director 2/128 cards in one chassis:

1. Connect to the switch and log in as admin.
2. Use the `configupload` command to back up the configuration of sw0 (slots 1 through 4).
3. Issue the `switchshutdown` command to ensure a graceful shutdown of sw0. Wait until the command finishes and displays the message:

Cleaning up kernel modules . . . . .Done

The following is a sample output from the command:

```
SW0:admin> switchshutdown
Stopping all switch daemons...Done.
Powering off slot 1...Done.
Powering off slot 4...Done.
Checking all slots are powered off...Done.
Cleaning up kernel modules.....Done
SW0:admin>
```

4. Shut down the power to the switch.

For details on the `switchshutdown` command, refer to the *HP StorageWorks Fabric OS 4.x command reference guide* or to the online help. For details on shutdown procedures, refer to the *HP StorageWorks Core Switch 2/64 and SAN Director 2/128 installation guide*.

5. Remove the Core Switch 2/64 CP cards from slots 5 and 6 of the chassis.
6. Insert the SAN Director 2/128 CP cards into slots 5 and 6 of the chassis.
7. Insert the SAN Director 2/128 port cards into the empty side of the chassis (slots 7 through 10).
8. Restore power to the switch.

By default, the switch starts up in single domain mode (one 128-port switch) with slots 1 through 4 set to faulty.

9. Connect to the switch and log in as admin.

**10.** Issue the slotshow command to view the status of the cards in each slot. The Core Switch 2/64 cards (ID = 2) show FAULTY status. For example:

```
slotshow
```

Slot	Blade Type	ID	Status
1	SW BLADE	2	FAULTY (9)
2	SW BLADE	2	FAULTY (9)
3	SW BLADE	2	FAULTY (9)
4	SW BLADE	2	FAULTY (9)
5	CP BLADE	5	ENABLED
6	CP BLADE	5	ENABLED
7	SW BLADE	4	ENABLED
8	SW BLADE	4	ENABLED
9	SW BLADE	4	ENABLED
10	SW BLADE	4	ENABLED

Issue the chassisconfig command to configure two domains. Use the -f option to suppress prompting for uploading the configuration and the 4 option to specify two 64-port switches (Blade ID 2 on slots 1–4, ID 4 on slots 7–10).

**11.** This command reboots the system.

```
chassisconfig -f 4
Current Option changed to 4
Restoring switch 0 configuration to factory defaults...
All account passwords have been successfully set to factory default.
Restoring switch 1 configuration to factory defaults...
All account passwords have been successfully set to factory default.
```

**12.** After the system reboots, log in again as admin to each logical switch.

Passwords have been changed to the defaults. You can either change the account passwords or press **Ctrl+c** to bypass prompts.

**13.** Issue the chassisconfig command without options to verify the change to two domains.

For example:

```
chassisconfig
Current Option: 4
```

**14.** Issue the slotshow command to verify that there are no faulty cards. If POST diagnostics are running, allow them to finish, which takes several minutes.

```
slotshow
```

Slot	Blade Type	ID	Status
1	SW BLADE	2	DIAG RUNNING POST1
2	SW BLADE	2	DIAG RUNNING POST1
3	SW BLADE	2	DIAG RUNNING POST1
4	SW BLADE	2	DIAG RUNNING POST1
5	CP BLADE	5	DIAG RUNNING POST1
6	CP BLADE	5	DIAG RUNNING POST1
7	SW BLADE	4	DIAG RUNNING POST1
8	SW BLADE	4	DIAG RUNNING POST1
9	SW BLADE	4	DIAG RUNNING POST1
10	SW BLADE	4	DIAG RUNNING POST1

**15.** Reissue the `slotshow` command until you see that POST diagnostics are finished and the status of all cards is Enabled. For example:

```
slotshow
```

Slot	Blade Type	ID	Status
1	SW BLADE	2	ENABLED
2	SW BLADE	2	ENABLED
3	SW BLADE	2	ENABLED
4	SW BLADE	2	ENABLED
5	CP BLADE	5	ENABLED
6	CP BLADE	5	ENABLED
7	SW BLADE	4	ENABLED
8	SW BLADE	4	ENABLED
9	SW BLADE	4	ENABLED
10	SW BLADE	4	ENABLED

**16.** Issue the `switchshow` command to verify that port initialization is complete (no ports are shown as Testing and all E\_Ports, F\_Ports, and L\_Ports are Online).

**17.** Use the `configdownload` command to restore the configuration of sw0 (saved in [step 2](#)).

**18.** Manually configure sw1 as desired.

---

## Setting the card beacon mode

When beaconing mode is enabled, the port LEDs flash amber in a running pattern from port 0 through port 15 and back again. The pattern continues until you turn it off. This can be used to locate a particular card.

To set the card beacon mode on:

1. Connect to the switch and log in as admin.
2. Issue the `bladebeacon` command with the following syntax at the command line:

```
bladebeacon slotnumber, mode
```

where `slotnumber` is the card where you want to enable beacon mode; this slot number must exist on the logical switch. 1 turns beaconing mode on, or 0 turns beaconing mode off.

Example:

```
switch:admin> bladebeacon 3, 1
switch:admin>
```

# 6 Routing traffic

This chapter contains procedures for configuring HP StorageWorks switch routing features. For details on the commands used in the procedures, refer to the *HP StorageWorks Fabric OS 4.x command reference guide*.

This chapter contains the following sections:

- [About routing policies](#), page 99
- [Specifying the routing policy](#), page 99
- [Assigning a static route](#), page 100
- [Specifying frame order delivery](#), page 100
- [Using dynamic load sharing](#), page 101
- [Viewing routing path information](#), page 102
- [Viewing routing information along a path](#), page 104

## About routing policies

All HP StorageWorks switches support *port-based* routing, in which the routing path chosen for an incoming frame is based only on the incoming port and the destination domain. To optimize port-based routing, enable the Dynamic Load Sharing feature (DLS) to balance the load across the available output ports within a domain.

The SAN Switch 4/32 allows you to tune routing performance with these additional routing policies:

- *Device-based* routing, in which the choice of routing path is based on the Fibre Channel addresses of the source device (SID) and the destination device (DID), improving path utilization for better performance
- *Exchange-based* routing, in which the choice of routing path is based on the SID, DID, and Fibre Channel originator exchange ID (OXID), optimizing path utilization for the best performance

Device-based and exchange-based routing require the use of DLS; when these policies are in effect, you cannot disable the DLS feature.

Using port-based routing, you can assign a *static route*, in which the path chosen for traffic never changes. In contrast, device-based and exchange-based routing policies always employ *dynamic path selection*.

## Specifying the routing policy

In addition to port-based routing, which all HP StorageWorks switches support, the SAN Switch 4/32 supports additional routing policies and allows you to specify the active routing policy using the `aptpolicy` command.

The following routing policies are supported:

- 1: Port-based path selection, which is the default on the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128

- 2: Device-based path selection on the SAN Switch 4/32 only
- 3: Exchange-based path selection, which is the default on the SAN Switch 4/32 only

The default policy usually provides the best performance. You should change the policy only if there is a performance problem that you cannot resolve in other ways.

You must disable the switch before changing the routing policy, and reenable it afterward.

In this example, the routing policy is changed from exchange-based to device-based:

```
switch:admin> aptpolicy
Current Policy: 3

3: Default Policy
1: Port Based Routing Policy
2: Device Based Routing Policy
3: Exchange Based Routing Policy
switch:admin> switchdisable
switch:admin> aptpolicy 2
Policy updated successfully.
switch:admin> switchenable
switch:admin> aptpolicy
Current Policy: 2
```

## Assigning a static route

Assign a static route only when the active routing policy is port-based. When device-based or exchange-based routing is active you cannot assign static routes.

To assign a static route, use the **urouteconfig** command. To remove a static route, use the **urouteremove** command.



**NOTE:** For the SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128, when you issue the **urouteconfig** command, two similar warning messages may be displayed if a platform conflict condition occurs. The first message appears when the static routing feature detects the condition. The second message appears when the dynamic load sharing feature detects the condition as it tries to rebalance the route.

A platform conflict occurs if a static route was configured with a destination port that is currently down. The static route is ignored in this case, in favor of a normal dynamic route. When the configured destination port comes back up, the system attempts to reestablish the static route, and the conflict can occur then.

## Specifying frame order delivery

In a stable fabric, frames are always delivered in order, even when the traffic between switches is shared among multiple paths. However, when topology changes occur in the fabric (for example, if a link goes down), traffic is rerouted around the failure, and some frames could be delivered out of order. Most destination devices tolerate out-of-order delivery, but some do not.

By default, out of order frame-based delivery is allowed to improve speed. You should force in-order frame delivery across topology changes only if the fabric contains destination devices that cannot tolerate occasional out-of-order frame delivery.

To force in-order frame delivery across topology changes:

1. Connect to the switch and log in as admin.
2. Issue the `iodset` command.



**NOTE:** This command can cause a delay in the establishment of a new path when a topology change occurs, and should be used with care.

To restore out-of-order frame delivery across topology changes:

1. Connect to the switch and log in as admin.
2. Issue the `iodreset` command.

## Using dynamic load sharing

The device-based and exchange-based routing policies depend on the Fabric OS dynamic load sharing feature (DLS) for dynamic routing path selection. When these policies are in force, DLS must be enabled and cannot be disabled.

When the port-based policy is in force, you can enable DLS to optimize routing. When DLS is enabled, it shares traffic among multiple equivalent paths between switches. DLS recomputes load sharing when a switch boots up, or each time an E\_Port or Fx\_Port goes online or offline.

To check and set DLS:

1. Connect to the switch and log in as admin.
2. Issue the `dlsshow` command to view the current DLS setting.

One of the following messages appears:

- DLS is set means that dynamic load sharing is turned on.
- DLS is not set means that dynamic load sharing is turned off.

3. Issue the `dlsset` command to enable DLS or issue the `dlsreset` command to disable it.

You cannot disable DLS when device-based or exchange-based routing policies are in effect.

Example:

```
switch:admin> dlsshow
DLS is not set
switch:admin> dlsset
switch:admin> dlsshow
DLS is set
switch:admin> dlsreset
switch:admin> dlsshow
DLS is not set
```

# Viewing routing path information

The `topologyshow` and `urouteshow` commands provide information about the routing path.

1. Connect to the switch and log in as admin.
2. Issue the `topologyshow` command to display the fabric topology as it appears to the local switch.

The following entries appear:

- Local Domain ID      Domain number of local switch
- Domain                Domain number of destination switch
- Metric                Cost of reaching destination domain
- Name                  The name of the destination switch
- Path Count           The number of currently active paths to the destination domain
- Hops                  The maximum number of hops to reach destination domain
- Out Port             The port that the incoming frame is forwarded to, in order to reach the destination domain
- In Ports             Input ports that use the corresponding Out Port to reach the destination domain
- Total Bandwidth      The maximum bandwidth of the out port
- Bandwidth Demand    The maximum bandwidth demand of the In Ports
- Flags                Always `D`, indicating a dynamic path

Example:

```
switch:admin> topologyshow
2 domains in the fabric; Local Domain ID: 1
Domain: 6
Metric: 500
Name: switch
Path Count: 4
Hops: 1
Out Port: 60
In Ports: None
Total Bandwidth: 2 Gbps
Bandwidth Demand: 0%
Flags: D
Hops: 1
Out Port: 61
In Ports: None
Total Bandwidth: 2 Gbps
Bandwidth Demand: 0%
Flags: D
Hops: 1
Out Port: 62
In Ports: None
Total Bandwidth: 2 Gbps
Bandwidth Demand: 0%
Flags: D
Hops: 1
Out Port: 58
In Ports: None
Total Bandwidth: 2 Gbps
Bandwidth Demand: 0%
Flags: D
```

**3.** Issue the **urouteshow** command to display unicast routing information.

For the HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32, use the following syntax:

```
urouteshow [portnumber] [, domainnumber]
```

For the Core Switch 2/64 and the SAN Director 2/128, use the following syntax:

```
urouteshow [slot/] [portnumber] [, domainnumber]
```

The following entries appear:

- Local Domain              Domain number of local switch
- In Ports                  Port from which a frame is received
- Domain                   Destination domain of incoming frame
- Out Port                  The port that incoming frame is forwarded to, in order to reach the destination domain
- Metric                   Cost of reaching destination domain
- Hops                     The maximum number of hops to reach destination domain
- Flags                    Indicates whether route is dynamic (D) or static (S). A static route is assigned using the command **urouteconfig**
- Next (Dom, Port)        Domain and port number of the next hop. These are the domain number and the port

The following example displays the routing information of all the active ports:

```
switch:admin> urouteshow
Local Domain ID: 3
In PortDomain Out Port Metric Hops Flags Next (Dom, Port)
-----
0 1 11 1000 1 D 1,0
11 2 0 1500 2 D 4,0
        4 16 500 1 D 4,0
16 1 27 1000 1 D 1,1
27 2 16 1500 2 D 4,16
4 0 29 500 1 D 4,0
```

The following example displays the routing information for port 11 on slot 1:

```
switch:admin> urouteshow 1/11
Local Domain ID: 3
In PortDomain Out Port Metric Hops Flags Next (Dom, Port)
-----
11 2 0 1500 2 D 4,0
        4 16 500 1 D 4,0
```

The following example displays the routing information of port 11 to domain 4 only:

```
switch:admin> urouteshow 1/11, 4
Local Domain ID: 3
In PortDomain Out Port Metric Hops Flags Next (Dom, Port)
-----
11 4 16 500 1 D 4,0
```

## Viewing routing information along a path

You can display detailed routing information from a source port (or area) on the local switch to a destination port (or area) on another switch. This routing information describes the full path that a data stream travels between these ports, including all intermediate switches.

1. Connect to the switch and log in as admin.
2. Issue the **pathinfo** command. In interactive mode, you can specify the following parameters for display:
  - Max hops              The maximum number of hops that the pathinfo frame is allowed to traverse
  - Domain                The destination domain ID
  - Source Port           The port number (or area number for the Core Switch 2/64 and the SAN Director 2/128) on which the switch receives frames
  - Destination Port     The output port that the frames use to reach the next hop on this path. For the last hop, the destination port
  - Basic stats           Basic statistics on every link

- Extended stats      Detailed statistics on every link
- Trace reverse path    The path from the destination switch back to the source switches
- Source route        The route that the frame is forced to follow a specified path to reach the destination
- Timeout              The maximum time to wait for a response from `pathinfo`, in seconds

Paths always originate on the local switch. The path destination can be specified by domain or port. By default, the path is the path taken by traffic from the source to destination port, but you can also specify all or portions of a path.

Refer to the *HP StorageWorks Fabric OS 4.x command reference guide* for details on this command.

The following example is from a SAN Switch 2/32 (other models provide similar information):

```

switch:admin> pathinfo

      Max hops: (1..127) [25]
      Domain: (1..239) [-1] 1
      Source port: (0..255) [-1]
      Destination port: (0..255) [-1]
      Basic stats (yes, y, no, n): [no]
      Extended stats (yes, y, no, n): [no]
      Trace reverse path (yes, y, no, n): [no]
      Source route (yes, y, no, n): [no]
      Timeout: (1..30) [10]

      Target port is Embedded

      Hop  In Port  Domain ID (Name)          Out Port    BW   Cost
      --  --  --
      0       E        10 (SW3900)           15        2G   500
      1       7        1 (swd3900TechPu)      E        -
switch:admin>

```

The information that `pathinfo` provides is:

- Hop      The hop number. The local switch is hop 0.
- In Port    The port that the frames come in from on this path. For hop 0, the source port.
- Domain ID The domain ID of the switch.
- Name      The name of the switch.
- Out Port   The output port that the frames use to reach the next hop on this path. For the last hop, the destination port.
- BW        The bandwidth of the output ISL, in Gbps. It does not apply to the embedded port.
- Cost      The cost of the ISL used by FSPF routing protocol. It applies only to an E\_Port.



# 7 Administering extended fabrics

This chapter contains procedures for using the HP Extended Fabrics licensed feature, which extends the distance that interswitch links (ISLs) can reach. To use extended ISL modes, you must first install the Extended Fabrics license. For details on obtaining and installing licensed features, see “[Maintaining licensed features](#)” on page 25.

This chapter contains the following sections:

- [About extended link buffer allocation](#), page 107
- [Fabric considerations](#), page 108
- [Choosing an extended ISL mode](#), page 108
- [Configuring an extended ISL](#), page 109
- [Trunking over distance](#), page 111

## About extended link buffer allocation

As the distance between switches and the link speed increase, additional *buffer-to-buffer credits* are required to maintain maximum performance. The number of credits reserved for a port depends on the switch model and on the extended ISL mode for which it is configured.

### SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128

Each *port group* (called a *quad* on these models) contains four ports and shares a common pool of credits. Because the number of credits available for use within each port group is limited, configuring ports for extended links on these models may cause other ports to become disabled if there are not enough buffer credits available; for example:

- If two 2-Gbps ports in a group are configured for L1 mode, each is allocated enough buffer-to-buffer credits to cause the other two ports in the group to become disabled (see [Table 15](#)).
- A port connected to a device that is in loopback mode may become disabled for lack of buffers if another port in that group is set to L2 mode.

### SAN Switch 4/32

Each port group contains eight ports and buffer credits are shared among all ports on the switch. *Buffer-limited* port technology allows all ports to remain operational, even when extended links are in use.

A buffer-limited port can come online with fewer buffer credits allocated than its configuration specifies, allowing it to operate at a reduced bandwidth instead of being disabled for lack of buffers.

Buffer-limited operation is supported for the L0 and LD extended ISL modes only, and is persistent across reboots, switch disabling and enabling, and port disabling and enabling.

# Fabric considerations

Consider these items that affect the fabric when you configure extended ISLs:

- The extended link configuration mode, L2 can reach 100 km at a speed of 2Gbps between HP Fabric OS v4.x switches. However, it only supports a distance of up to 60 km if the link is established between HP Fabric OS v3.x and 4.x switches.
- The standard-distance and long-distance ISL modes cannot be enabled at the same time.
- Balance the number of long-distance ISL connections and core-to-edge ISL connections within a switch. Configuring long-distance ISLs between core and edge switches is possible, but is not a practice recommended by HP.
- Starting with Fabric OS v4.4.0, VC translation link initialization (an option of the `portcfglongdistance` command) is enabled by default for long-distance links. For previous Fabric OS versions that support this option, it was disabled by default. To avoid inconsistency in the fabric, make sure that this value is enabled on both ends of the link. To connect to switches running Fabric OS versions earlier than v4.0.2 and v3.0.2c, make sure that VC translation link initialization is disabled, because these versions do not support it.

## Choosing an extended ISL mode

**Table 15** lists the extended ISL modes, which you can configure with the `portcfglongdistance` command when the Extended Fabrics license is activated. For standard ISL modes that do not require Extended Fabrics licensing, see [Table 4](#) on page 31.

**Table 15** Extended ISL modes

Mode	Description	Maximum ISL distance (km)	Earliest Fabric OS release
L0.5	Level 0.5 static mode (designated LM when listed with the <code>portcfgshow</code> command).	25 km at 1, 2, or 4 Gbps	v3.1.0, v4.1.0
L1	Level 1 static mode.	50 km at 1, 2, or 4 Gbps	All
L2	Level 2 static mode.	64 or 100 km at 1, 2, or 4 Gbps <sup>1</sup>	All
LD	Dynamic mode uses automatic distance detection for a user-specified distance.	<ul style="list-style-type: none"><li>SAN Switch 4/32:<ul style="list-style-type: none"><li>500 km at 1 Gbps</li><li>250 km at 2 Gbps</li><li>100 km at 4 Gbps</li></ul></li><li>Other supported models<sup>1</sup>:<ul style="list-style-type: none"><li>100 or 200 km at 1 Gbps</li><li>64 or 100 km at 2 Gbps</li></ul></li></ul>	v3.1.0, v4.1.0, v 4.4.0 (depending on the model)

1. The maximum ISL distance depends on the version of HP StorageWorks ASIC installed in the switch.

The dynamic long-distance mode (LD) automatically configures the number of buffer credits required, based on the actual link distance.

For the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128, the number of ports that can be configured per port group at various distances is summarized in [Table 16](#).

**Table 16** SAN Switch 2/8V, 2/16V, 2/32, Core Switch 2/64, and SAN Director 2/128

Speed (Gbps)	Number of Ports Allowed at Distance (km)					
	10	25	50	100	250	500
1	4	4	L1 mode: 4 LD mode: 3	up to 2	n/a	n/a
2	4	3	n/a	1	n/a	n/a
4	n/a	n/a	n/a	n/a	n/a	n/a

For the SAN Switch 4/32, the number of ports that can be configured at various distances is summarized in [Table 17](#).

**Table 17** SAN Switch 4/32

Speed (Gbps)	Number of Ports Allowed at Distance (km)					
	10	25	50	100	250	500
1	32	32	32	up to 15	up to 6	up to 3
2	32	32	15	up to 7	up to 3	n.a.
4	32	15	7	up to 3	n.a.	n.a.

## Configuring an extended ISL

Before configuring the ISL, ensure that the following conditions are met:

- Extended ISL support for HP StorageWorks 1 GB switches is limited as follows:
  - Extended ISLs are not supported on HP StorageWorks 1 GB switches that are part of a fabric that contains other HP StorageWorks switches.
  - To support extended ISLs between HP StorageWorks 1 GB switches, the `fabric.ops.mode.longDistance` parameter to 1 on all switches in the fabric. Each switch must be disabled before setting this parameter. The SAN Switch 4/32 cannot be part of such a fabric.
  - For fabrics that contain a mix of HP StorageWorks switches, the `fabric.ops.mode.longDistance` parameter must be set to 0 (the default). (Under certain circumstances this mode needs to be enabled on switches running Fabric OS v3.x or v4.x. Talk to your switch provider for details.)
  - The ports on both ends of the ISL must have the same configuration.
  - Use only qualified SFPs. Refer to your model's hardware manual for details on SFPs.

To configure an extended ISL:

1. Connect to the switch and log in as admin.
2. If the fabric contains a mix of switches, use the `configure` command to make sure the fabric-wide configuration parameter `fabric.ops.mode.longDistance` is set to 0 on all switches in the fabric.

If the fabric contains only HP StorageWorks 1 GB switches, use the `switchDisable` command to disable the switch and then use the `configure` command to set the fabric-wide configuration parameter `fabric.ops.mode.longDistance` to 1 on all switches in the fabric. The SAN Switch 4/32 cannot be part of such a fabric.

3. Issue the `portcfglongdistance` command, using the following syntax:

```
portcfglongdistance [slot/]port [distance_level] [vc_translation_link_init]  
[desired_distance]
```

where:

*slot* Specifies the slot number for the Core Switch 2/64 and the SAN Director 2/128. This option is not applicable to fixed-port switches. The slot number must be followed by a slash ( / ) and the port number.

*port* Specifies the port number.

*distance\_level* Specifies the ISL mode to be set on the port (see [Table 15](#) on page 108).

*vc\_translation\_link\_init* Is the extended link initialization sequence, which is an enhanced link reset protocol that avoids excessive resetting of ports.

By default this option is set to 1 (enabled).

To prevent fabric segmentation, this option must be set to 0 (disabled) when connecting to switches running Fabric OS versions earlier than v3.0.2c or v4.0.2.

It must be set to 1 (enabled) in the following circumstances:

- When configuring a trunk over extended fabrics
- For optimal performance of ISLs between switches running Fabric OS v3.0.2 or later, or Fabric OS v4.0.2 or later

*desired\_distance* Required for a port configured for LD mode. Specifies the desired distance, in kilometers, for the link. The specified value is the upper limit for calculating buffer availability for the port. If the measured distance is more than the specified *desired\_distance*, the port is allocated the number of buffers required by the specified desired distance. (Fabric OS versions earlier than v4.4.0 do not support this parameter.)

4. Repeat [step 3](#) for the remote extended ISL port. Both the local and remote extended ISL ports must be configured to the same distance level. When the connection is initiated, the fabric reconfigures.

The following example configures slot 1, port 1 for the LD link distance mode, enables the extended link initialization sequence, and sets the desired distance to 50 kilometers:

```
switch:admin> portcfglongdistance 1/1 LD 1 50  
switch:admin>
```



**NOTE:** In rare cases, reconfiguring a port to LD from one of the other modes can result in the port being disabled for lack of buffers. If this occurs:

In Fabric OS v4.2.x, reconfigure the disabled LD port back to the original mode.

In Fabric OS v4.4.0 and later, specify a slightly shorter distance for the *desired\_distance* parameter in the `portcfglongdistance` command.

---

## Trunking over distance

See "Trunking over extended fabrics" on page 119.



# 8 Administering ISL trunking

This chapter contains procedures for using the HP ISL Trunking licensed feature, which optimizes the use of bandwidth by allowing a group of interswitch links to merge into a single logical link.

This chapter contains the following sections:

- [Standard trunking criteria](#), page 113
- [Fabric considerations](#), page 114
- [Initializing trunking on ports](#), page 115
- [Monitoring traffic](#), page 115
- [Enabling and disabling ISL trunking](#), page 116
- [Setting port speeds](#), page 117
- [Displaying trunking information](#), page 119
- [Trunking over extended fabrics](#), page 119
- [Troubleshooting trunking problems](#), page 120

To use trunking, you must first install the ISL Trunking license. For details on obtaining and installing licensed features, see “[Maintaining licensed features](#)” on page 25.

For detailed information about trunking commands, refer to online help or the *HP StorageWorks Fabric OS 4.x command reference guide*.

## Standard trunking criteria

Observe the following criteria for standard distance trunking:

- There must be a direct connection between participating switches.
- Trunk ports must reside in the same port group.
- Trunk ports must run at the same speed (either 2 Gbps or 4 Gbps).
- Trunk ports must be set to the same ISL mode (LO or LE). For details on these modes, see [Table 4](#) on page 31.
- Trunk ports must be E\_Ports.
- Cable lengths for participating links should differ by no more than 30 meters.
- The `switch.interopMode` parameter must be set to 0. See “[Configuring interoperability mode](#)” on page 219 for information and procedures related to interoperability mode.
- The port ISL mode must be disabled (using the `portIslMode` command).

# Fabric considerations

The ISL Trunking feature is provided with the Fabric OS and can be activated by entering a license key, which is available from the switch supplier. When the ISL Trunking license is activated, trunking is automatically implemented for any eligible ISLs.

A license must be activated on each switch that participates in trunking. For the Core Switch 2/64, a single license key enables the feature on both logical switches.

To use ISL Trunking in the fabric, the fabric must be designed to allow trunking groups to form. To identify the most useful trunking groups, evaluate the traffic patterns before designing/redesigning the fabric. This also applies to the SAN Director 2/128 configured with two domains.

ISL Trunking can be used to simplify SAN design and improve performance. When designing the SAN, consider the following recommendations in addition to the standard guidelines for SAN design:

- Evaluate the traffic patterns within the fabric.
- Place trunking-capable switches adjacent to each other.

This maximizes the number of trunking groups that can form. If you are using a core/edge topology, place trunking-capable switches at the core of the fabric and any switches that are not trunking-capable at the edge of the fabric.

- Activate an ISL Trunking license on each switch that is to participate in a trunking group.
- Cable lengths for participating links should differ by no more than 30 meters.
- When connecting two switches with two or more ISLs, ensure that all trunking requirements are met to allow a trunking group to form.
- Determine the optimal number of trunking groups between each set of linked switches, depending on traffic patterns and port availability.

The goal is to avoid traffic congestion without unnecessarily using ports that could be used to attach other switches or devices. Consider these points:

- Each physical ISL uses two ports that could otherwise be used to attach node devices or other switches.
- Trunking groups can be used to resolve ISL oversubscription if the total capability of the trunking group is not exceeded.
- Consider how the addition of a new path affects existing traffic patterns:
  - A trunking group has the same link cost as the master ISL of the group, regardless of the number of ISLs in the group. This allows slave ISLs to be added or removed without causing data to be rerouted, because the link cost remains constant.
  - The addition of a path that is shorter than existing paths causes traffic to be rerouted through that path.
  - The addition of a path that is longer than existing paths may not be useful because the traffic chooses the shorter paths first.
- Plan for future bandwidth addition to accommodate increased traffic.

For trunking groups over which traffic is likely to increase as business requirements grow, consider leaving one or two ports in the group available for future nondisruptive addition of bandwidth.

- Consider creating redundant trunking groups where additional ports are available or paths are particularly critical.  
This helps to protect against oversubscription of trunking groups, multiple ISL failures in the same group, and the rare occurrence of an ASIC failure.
- To provide the highest level of reliability, deploy trunking groups in redundant fabrics to further ensure ISL failures do not disrupt business operations.

---

## Initializing trunking on ports

After you unlock the ISL Trunking license, you must reinitialize the ports being used for ISLs so that they recognize that trunking is enabled. This procedure needs to be performed only once.

To reinitialize the ports, you can either disable and then reenable the switch, or disable and then reenable the affected ports.

To disable and reenable the switch:

1. Connect to the switch and log in as admin.
2. Issue the `switchdisable` command.
3. Issue the `switchenable` command.

To disable and reenable ports:

1. Connect to the switch and log in as admin.
2. Issue the `portdisable` command. The format is:

```
portdisable [slot/]port
```

*slot* is the slot number (Core Switch 2/64 and SAN Director 2/128 only) and *port* is the port number of the port you want to disable.

3. Issue the `portenable` command. The format is:

```
portenable [slot/]port
```

*slot* is the slot number (Core Switch 2/64 and SAN Director 2/128 only) and *port* is the port number of the port you want to enable.

---

## Monitoring traffic

To implement ISL Trunking effectively, you must monitor fabric traffic to identify congested paths or to identify frequently dropped links. While monitoring changes in traffic patterns, you can adjust the fabric design accordingly, such as by adding, removing, or reconfiguring ISLs and trunking groups in problem areas.

There are three methods of monitoring fabric traffic:

- HP Advanced Performance Monitoring monitors traffic flow and allows you to view the impact of different fabric configurations on performance. See “[Administering advanced performance monitoring](#)” on page 143 for additional information.
- HP Fabric Watch allows you to monitor traffic flow through specified ports on the switch and send alerts when the traffic exceeds or drops below configurable thresholds. Refer to the *HP StorageWorks Fabric OS 4.x Fabric Watch user guide* for additional information.

- Use the **portperfshow** command as described in the following procedure to record traffic volume for each port in your fabric over time.

To use the **portperfshow** command:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
portperfshow [interval]
```

where *interval* is the number of seconds between each data-gathering sample (the default is one sample every second).

3. Record the traffic flow for each port participating in an ISL.
4. Repeat step 1 through step 3 for each switch in the fabric until all ISL traffic flow is captured.

In a large fabric, it may be necessary to only identify and capture the key ISLs. However, you may want to continue this process throughout the day (or an entire work cycle), to capture varying traffic patterns under different conditions.

The following example shows a switch without trunking, and indicates that ports 0 through 2 are underutilized and ports 4 and 5 are congested:

```
switch:admin> portperfshow
      0      1      2      3    4567 Total
-----
      0      0      0   145m204m202m0168m  719
      0      0      0   145m206m208m0186m  745
switch:admin>
```

The following example shows traffic flowing through a trunking group (ports 5, 6, and 7). After port 6 fails, traffic is redistributed over the remaining two links in the group, ports 5 and 7:

```
switch:admin> portperfshow
      0      1      2      3    4567 Total
-----
      0      0      0      0  0145m144m145m  434
      0      0      0      0  0144m143m144m  431
      0      0      0      0  0162m0162m   324
      0      0      0      0  0186m0186m   372
      0      0      0      0  0193m0192m   385
      0      0      0      0  0202m0202m   404
      0      0      0      0  0209m0209m   418
switch:admin>
```

## Enabling and disabling ISL trunking

You can enable or disable ISL Trunking for one port or for an entire switch. To enable or disable ISL Trunking on one port:

1. Connect to the switch and log in as admin.

- Issue the **portcfgtrunkport** command. The format is:

```
portcfgtrunkport slotnumber/portnumber 1|0
```

where:

*slotnumber* Specifies the number of the slot in which the port card containing the port is located (this operand is required only for switches with slots, such as the Core Switch 2/64 and SAN Director 2/128).

*portnumber* Specifies the number of the port on which you want to enable or disable trunking.

*1|0* Enables (1) or disables (0) trunking on the specified port.

The following example enables trunking on slot 1, port 3:

```
switch:admin> portcfgtrunkport 1/3 1
done.
switch:admin>
```

To enable or disable ISL Trunking for all of the ports on a switch:

- Connect to the switch and log in as admin.
- Issue the **switchcfgtrunk** command. The format is:

```
switchcfgtrunk 1|0
```

where 1 enables and 0 disables ISL Trunking for all ports on the switch.

The following example enables trunking all ports in the switch.

```
switch:admin> switchcfgtrunk 1
Committing configuration...done.
switch:admin>
```

---

## Setting port speeds

You can set the port speed for one port or for an entire switch. Trunked ports must be set to the same speed.

To set the speed for one port:

- Connect to the switch and log in as admin.

2. Issue the `portcfgspeed` command. The format is:

```
portcfgspeed slotnumber/portnumber speedlevel
```

where:

*slotnumber* Specifies the switch slot (this operand is required only for switches with slots, such as the Core Switch 2/64 and SAN Director 2/128).

*portnumber* Specifies the port number.

*speedlevel* Specifies the speed of the link:

- 0 is auto-negotiating mode. The port configures for the highest speed.
- 1 is 1 Gbps mode, which fixes the port at a speed of 1 Gbps. Changing the speed to 1Gbps causes the port to be excluded from the trunk group.
- 2 is 2 Gbps mode, which fixes the port at a speed of 2 Gbps.
- 4 is 4 Gbps mode, which fixes the port at a speed of 4 Gbps. (SAN Switch 4/32 only.)

The following example sets the speed for port 3 on slot 2 to 2 Gbps:

```
switch:admin> portcfgspeed 2/3 2
done.
switch:admin>
```

The following example sets the speed for port 3 on slot 2 to auto-negotiate:

```
switch:admin> portcfgspeed 2/3 0
done.
switch:admin>
```

To set the speed for all ports on the switch:

1. Connect to the switch and log in as admin.
2. Issue the `switchcfgspeed` command. The format is:

```
switchcfgspeed speedlevel
```

where:

*speedlevel* Specifies the speed of the link:

- 0 is auto-negotiating mode. The port configures for the highest speed.
- 1 is 1 Gbps, which changes the speed for all of the ports to 1 Gbps, (eliminating any ISL Trunking group).
- 2 is 2 Gbps mode, which fixes the port at a speed of 2 Gbps.
- 4 is 4 Gbps mode, which fixes the port at a speed of 4 Gbps. (SAN Switch 4/32 only.)

The following example sets the speed for all ports on the switch to 2 Gbps:

```
switch:admin> switchcfgspeed 2
Committing configuration...done.
switch:admin>
```

The following example sets the speed for all ports on the switch to auto-negotiate:

```
switch:admin> switchcfgspeed 0
Committing configuration...done.
switch:admin>
```

---

## Displaying trunking information

Use the `trunkshow` command to display the following information about ISL Trunking groups:

- Number identifier.
- Port-to-port connections, listed in the format *local port number -> remote port number*.
- WWNs of the remote switches.
- Deskew values (the time difference, in nanoseconds divided by 10, for traffic to travel over each ISL as compared to the shortest ISL in the group). The system automatically sets the minimum deskew value of the shortest ISL to 15.
- Master ports.

To display trunking information:

1. Connect to the switch and log in as admin.
2. Issue the `trunkshow` command.

The following example shows three trunking groups (1, 2, and 3); ports 1, 4, and 14 are masters:

```
switch:admin> trunkshow
1: 1 -> 1 10:00:00:60:69:04:10:83  deskew 16 Master
          0 -> 0 10:00:00:60:69:04:10:83  deskew 15
2: 4 -> 4 10:00:00:60:69:04:01:94  deskew 16 Master
          5 -> 5 10:00:00:60:69:04:01:94  deskew 15
          7 -> 7 10:00:00:60:69:04:01:94  deskew 17
          6 -> 6 10:00:00:60:69:04:01:94  deskew 16
3:14 -> 14 10:00:00:60:69:04:10:83  deskew 16 Master
          15 -> 15 10:00:00:60:69:04:10:83  deskew 15
switch:admin>
```

---

## Trunking over extended fabrics

In addition to the criteria listed in “[Standard trunking criteria](#)” on page 113, observe the following criteria for trunking over extended fabrics:

- ISL Trunking over extended fabrics is supported on switches running Fabric OS v3.2.0 (or later) or v4.4.x (or later).
- Extended Fabrics and ISL Trunking licenses are required on all participating switches.
- The `vc_translation_link_init` parameter must be set the same on all ports in an extended trunk. (For details on this parameter, see [page 110](#).)

# Troubleshooting trunking problems

If you have difficulty with trunking, try the solutions in this section.

## Listing link characteristics

If a link that is part of an ISL Trunk fails, use the `trunkdebug` command to troubleshoot the problem, as shown in the following procedure:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
trunkdebug port, port
```

where `port` specifies the number of a port in an ISL Trunking group.

The `trunkdebug` command displays the possible reason that two ports cannot be trunked, including the following reasons:

- The switch does not support trunking.
- A trunking license is required.
- Trunking is not supported in switch interoperability mode.
- Port trunking is disabled.
- The port is not an E\_Port.
- The port is not 2 Gbps or 4 Gbps.
- The port connects to different switches.
- The ports are not same speed, or they are not set to a valid speed.
- The ports are not set to the same long-distance mode.
- Local or remote ports are not in same port group.
- The difference in the cable length among trunked links is greater than the allowed difference.

This example shows that port 3 is not configured as an E\_Port:

```
switch:admin> trunkdebug 3 5
port 3 is not E port
switch:admin>
```

## Recognizing buffer underallocation

For the SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128, if there is an underallocation or overcommitment of buffers to ports configured for extended trunking, the switches at both ends of the trunk try to disable some ports so that others can operate using the available buffers. (Standard trunks are not affected by buffer allocation.) This issue does not apply to the SAN Switch 4/32.

A port disabled at one end causes all the disabled ports at the other end to become enabled. Some of these enabled ports become disabled due to a lack of buffers, which in turn triggers ports to be enabled once again at the other end. While the system is stabilizing the buffer allocation, it warns that ports are disabled due to lack of buffers, but it does not send a message to the console when buffers are enabled. The system requires a few passes to stabilize the buffer allocation. Ultimately, the number of ports for

which buffers are available come up and stabilize. You should wait for stabilization, and then proceed with correcting the buffer allocation situation.



# 9 Administering advanced zoning

This chapter provides procedures for using the HP Advanced Zoning feature and consists of the following sections:

- [Zoning terminology](#), page 123
- [Zoning concepts](#), page 124
- [Creating and managing zone aliases](#), page 131
- [Creating and maintaining zones](#), page 133
- [Creating and modifying zoning configurations](#), page 134
- [Managing zoning configurations in a fabric](#), page 137
- [Using zoning to administer security](#), page 139
- [Resolving zone conflicts](#), page 140

## Zoning terminology

The following Advanced Zoning terms are used in the procedures:

- A zone is a region within the fabric where a specified group of fabric-connected devices (called *zone members*) have access to one another. When zoning is enabled, objects not explicitly defined in a zone are isolated, and members in the zoned fabric do not have access to them.
- A group of one or more zones is called a *zone configuration*.
- The complete set of all zone members defined in a fabric is called the *defined zone configuration*.
- Zoning procedures change zone objects in the defined configuration. When you enable a configuration with the `cfgenable` command, it becomes the *effective zone configuration*.
- A copy of the defined zone configuration (plus the name of the effective zone configuration) can be saved with the `cfgsave` command. The resulting *saved zone configuration* is restored after a switch reboot. If you make changes to the defined zone configuration but do not save them, there will be differences between the defined zone configuration and the saved zone configuration.

To use zoning, you must install zoning licenses on all the switches in the fabric before attempting to bring a switch into the fabric. If a zoning license is removed, you must make sure it is reinstalled properly on the affected switch before attempting the `cfgenable` zoning operation. Failure to follow these steps can cause inconsistency of the zoning configuration on the affected switches if a zoning operation is attempted from a remote switch in the fabric. On the affected switches, an error message indicates that the zoning license is missing.

To list the commands associated with zoning, use the `zonehelp` command.

For detailed information on the zoning commands used in the procedures, refer to the *HP StorageWorks Fabric OS 4.x command reference guide* or to the online man page for each command.

# Zoning concepts

Before using the procedures, you should become familiar with the zoning concepts described in the following sections.

## Zone types

[Table 18](#) summarizes the types of zoning and [Table 19](#) summarizes the primary forms.

**Table 18** Types of zoning

Zone type	Description
Storage-based	Storage units typically implement LUN-based zoning, also called <i>LUN masking</i> . LUN-based zoning limits access to the LUNs on the storage port to the specific WWN of the server HBA. It is needed in most SANs. It functions during the probe portion of the SCSI initialization. The server probes the storage port for a list of available LUNs and their properties. The storage system compares the WWN of the requesting HBA to the defined zone list, and returns the LUNs assigned to the WWN. Other LUNs on the storage port are not made available to the server.
Host-based	Host-based zoning can implement WWN or LUN masking.
Fabric-based	Fabric switches implement fabric-based zoning, in which the zone members are identified by WWN or port location in the fabric. Fabric-based zoning is also called <i>name server-based</i> or <i>soft zoning</i> .  HP StorageWorks switches might also provide additional hardware enforcement of the zone. When a device queries the fabric Name Server, the Name Server determines the zones in which the device belongs. The server returns information on all members of the zones in the fabric to the device. Devices in the zone are identified by node WWN, port WWN, or domain, port of the switch to which the device is connected.  Fabric-based zoning is perhaps the most controversial aspect of zoning. There are several approaches for implementing fabric-based zoning; all of them work, in most cases. However, there are pros and cons to each form.

**Table 19** Approaches to fabric-based zoning

Zoned by	Description
Single HBA	Zoning by single HBA most closely re-creates the original SCSI bus. Each zone created has only one HBA (initiator) in the zone; each of the target devices is added to the zone. Typically, a zone is created for the HBA and the disk storage ports are added. If the HBA also accesses tape devices, a second zone is created with the HBA and associated tape devices in it. In the case of clustered systems, it could be appropriate to have an HBA from each of the cluster members included in the zone; this is equivalent to having a shared SCSI bus between the cluster members and presumes that the clustering software can manage access to the shared devices. In a large fabric, zoning by single HBA requires the creation of possibly hundreds of zones; however, each zone contains only a few members. Zone changes affect the smallest possible number of devices, minimizing the impact of an incorrect zone change. This zoning philosophy is the preferred method.
Application	Zoning by application typically requires zoning multiple, perhaps incompatible, operating systems into the same zones. This method of zoning creates the possibility that a minor server in the application suite could disrupt a major server (such as a web server) disrupting a data warehouse server. Zoning by application can also result in a zone with a large number of members, providing greater susceptibility to administrative errors, such as registered state change notifications (RSCNs) going out to a larger group than necessary.
Operating system	Zoning by operating system has issues similar to zoning by application. In a large site, this type of zone can become very large and complex. When zone changes are made, they typically involve applications rather than a particular server type. If members of different operating system clusters can see storage assigned to another cluster, they might attempt to own the other cluster's storage and compromise the stability of the clusters.
Port allocation	Avoid zoning by port allocation unless the administration team has very rigidly enforced processes for port and device allocation in the fabric. It does, however, provide some positive features. For instance, when a storage port, server HBA, or tape drive is replaced, the change of WWN for the new device is of no consequence. As long as the new device is connected to the original port, it continues to have the same access rights. The ports on the edge switches can be pre-associated to storage ports, and control of the fan-in ratio (the ratio of the input port to output port) can be established. With this pre-assigning technique, the administrative team cannot overload any one storage port by associating too many servers with it.
No fabric zoning	Using no fabric zoning is the least desirable zoning option because it allows devices to have unrestricted access on the fabric. Additionally, any device attached to the fabric, intentionally or maliciously, likewise has unrestricted access to the fabric. This form of zoning should be utilized only in a small and tightly controlled environment, such as when host-based zoning or LUN masking is deployed.

## Zone objects

A zone object is any device in a zone, such as the:

- Physical port number or area ID on the switch
- Node World Wide Name (N-WWN)
- Port World Wide Name (P-WWN)

Zone objects identified by port number or area number are specified as a pair of decimal numbers in the form  $d, \text{area}$  ( $d$  is the domain ID of the switch and  $\text{area}$  is the area number on that switch).

For example, on the Core Switch 2/64 and the SAN Director 2/128,  $4, 46$  specifies port 14 in slot number 3 (domain ID 4, area 46). On fixed-port models,  $3, 13$  specifies port 13 in switch domain ID 3.

When the physical port number specifies a zone object, then all devices connected to that port are in the zone. If the physical port is an arbitrated loop, then all devices on the loop are part of the zone.

World Wide Names are specified as 8-byte (16-digit) hexadecimal numbers, separated by colons: for example,  $10:00:00:90:69:00:00:8a$ . When a node name specifies a zone object, all ports on such a device are in the zone. When a port name specifies a zone object, only the single port is in the zone.

The types of zone objects used to define a zone can be mixed and matched. For example, a zone defined with the zone objects  $2,12; 2,14; 10:00:00:80:33:3f:aa:11$  contains the devices connected to domain 2, ports 12 and 14, and a device with the WWN (either node name or port name)  $10:00:00:80:33:3f:aa:11$  that is connected on the fabric.

## Zone aliases

A zone alias is a name assigned to a device or a group of devices. By creating an alias, you can assign a familiar name to a device or group multiple devices into a single name. This simplifies cumbersome data entry and allows an intuitive naming structure (such as using `NT_Hosts` to define all NT hosts in the fabric).

Zone aliases also simplify repetitive entry of zone objects, such as port numbers or a WWN. For example, you can use the name `Eng` as an alias for  $10:00:00:80:33:3f:aa:11$ .

A useful convention is to name zones for the initiator they contain. For example, if you use the alias `SRV_MAILSERVER_SLT5` to designate a mail server in PCI slot 5, then the alias for the associated zone is `ZNE_MAILSERVER_SLT5`. This clearly identifies the server host bus adapter (HBA) associated with the zone.

Zone configuration naming is more flexible. One configuration should be named `PROD_fabricname`, where `fabricname` is the name that the fabric has been designated. The purpose of the PROD configuration is to easily identify the configuration that can be implemented and provide the most generic services. If other configurations are used for specialized purposes, names such as `BACKUP_A`, `RECOVERY_2`, and `TEST_18jun02` can be used.

## Zone configurations

A zone configuration is a group of one or more zones. A zone can be included in more than one zone configuration. When a zone configuration is in effect, all zones that are members of that configuration are in effect.

The different types of zone configurations are:

- Defined Configuration, which is the complete set of all zone objects defined in the fabric.
- Effective Configuration, which is a single zone configuration that is currently in effect. The effective configuration is built when an administrator enables a specified zone configuration.

- Saved Configuration, which is a copy of the defined configuration plus the name of the effective configuration, which is saved in flash memory by the `cfgSave` command. (You can also use the `configupload` command to provide a backup of the zoning configuration and the `configdownload` command to restore the zoning configuration.) There might be differences between the saved configuration and the defined configuration if the system administrator has modified any of the zone definitions and has not saved the configuration.
- Disabled Configuration, which is the effective configuration is removed from flash memory.

On power-up, the switch automatically reloads the saved configuration. If a configuration was active when it was saved, the same configuration is reinstated with an autorun of the `cfgEnable` command.

You can establish a zone by identifying zone objects using one or more of the following zoning schemes:

- Domain, port number level. All members are specified by *domain ID, port number, or domain, area number* pair or aliases, described in "Zone aliases" on page 126.
- World Wide Name (WWN) level. All members are specified only by WWNs or aliases of WWNs. Members can be node or port versions of the WWN.
- Mixed zoning. A zone containing members specified by a combination of *domain, port number, and/or domain, area number* and WWN.

## Zoning enforcement

Software-enforced and hardware-enforced zoning are supported.

### Software-enforced zoning

Software-enforced zoning limits access to data by segmenting a fabric into virtual private SANs. Software-enforced zoning prevents hosts from discovering unauthorized target devices, while hardware-enforced zoning prevents a host from accessing a device it is not authorized to access.

Software-enforced zoning:

- Is also called *soft zoning*, *Name Server zoning*, *fabric-based zoning*, or *session-based zoning*.
- Is available on 1-Gbps, 2-Gbps, and 4-Gbps platforms.
- Prevents hosts from discovering unauthorized target devices.
- Ensures that the Name Server does not return any information to an unauthorized initiator in response to a Name Server query.
- Is always active whenever a zone configuration is in effect.
- Does not prohibit access to the device. If an initiator has knowledge of the network address of a target device, it does not need to query the Name Server to access it, which could lead to undesired access to a target device by unauthorized hosts.
- Is exclusively enforced through selective information presented to end nodes through the fabric Simple Name Server (SNS). When an initiator queries the Name Server for accessible devices in the fabric, the Name Server returns only those devices that are in the same zone as the initiator. Devices that are not part of the zone are not returned as accessible devices.

### Hardware-enforced zoning

Hardware-enforced zoning is specified without using the mixed zoning scheme. HP StorageWorks switches augment software-enforced zoning with hardware enforcement. The exact methodology varies on different switch models.

Hardware-enforced zoning (also called *hard zoning*):

- Prevents a host from accessing a device it is not authorized to access.

- Checks each frame before it is delivered to a zone member and discards it if there is a zone mismatch. When hardware-enforced zoning is active, the HP StorageWorks switch monitors the communications and blocks any frames that do not comply with the effective zone configuration. The switch performs this blocking at the transmit side of the port on which the destination device is located.
- Is enforced at the ASIC level. Each ASIC maintains a list of source port IDs that have permission to access any of the ports on that ASIC.

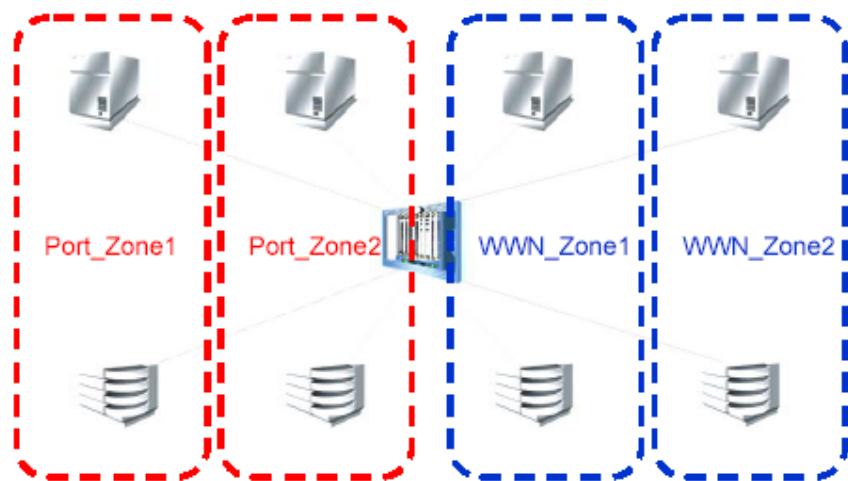
Fabric OS uses hardware-enforced zoning (on a per-zone basis) whenever the fabric membership or zone configuration changes.

**Table 20** shows the various HP StorageWorks switches, the hardware zoning methodology for each, and tips for best usage.

**Table 20** Enforcing hardware zoning

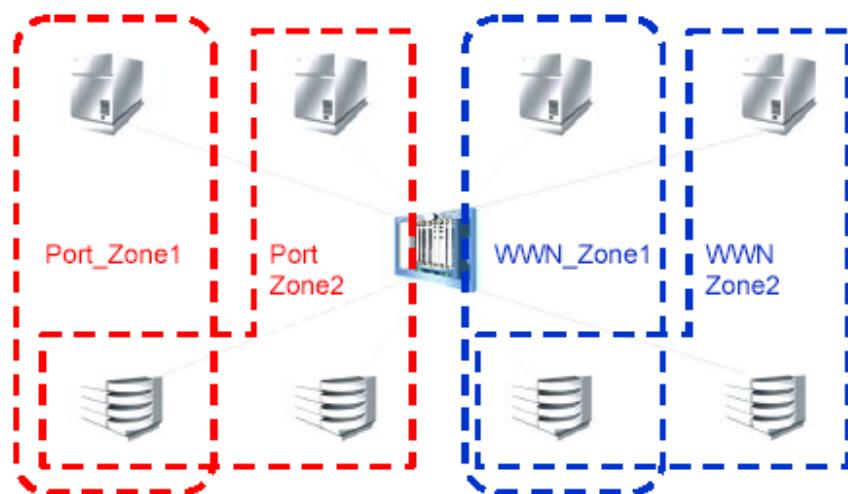
Fabric type	Methodology	Best practice
1-GB switches	<p>Enables hardware-enforced zoning only on domain, port zones; WWN or mixed zones are not hardware-enforced. Any domain, port zone that overlaps a mixed or WWN zone is not hardware-enforced.</p> <p>An overlap occurs when a member specified by WWN is connected to a port in a domain, port zone. The domain, port zone loses its hardware enforcement even though a review of the zone configuration does not indicate it.</p>	Use <i>domain, port</i> identifiers. Do not identify a zone member by its WWN.
2-GB switches, Core Switch 2/64, and SAN Director 2/128	<p>Enables hardware-enforced zoning on domain, port zones and WWN zones. Overlap of similar zone types does not result in the loss of hardware enforcement. Overlap with other zone type results in the loss of hardware enforcement.</p> <p>As in the 1-GB switches, connecting a device specified by WWN into a port specified in a domain, port zone results in loss of the hardware enforcement in both zones.</p>	Use either WWN or <i>domain, port</i> identifiers.
Mixed switches	<p>Enables hardware-enforced zoning according to each switch type. Use the <code>portzoneshow</code> command to find the switch type to which a device is attached.</p>	<p>Use <i>domain, port</i> identifiers.</p> <p>You can use WWN identifiers if you place disk and tape targets on Core Switch 2/64 and SAN Director 2/128, and do not use <i>domain, port</i> identifiers.</p>

[Figure 1](#) shows a fabric with four non-overlapping hardware-enforced zones.

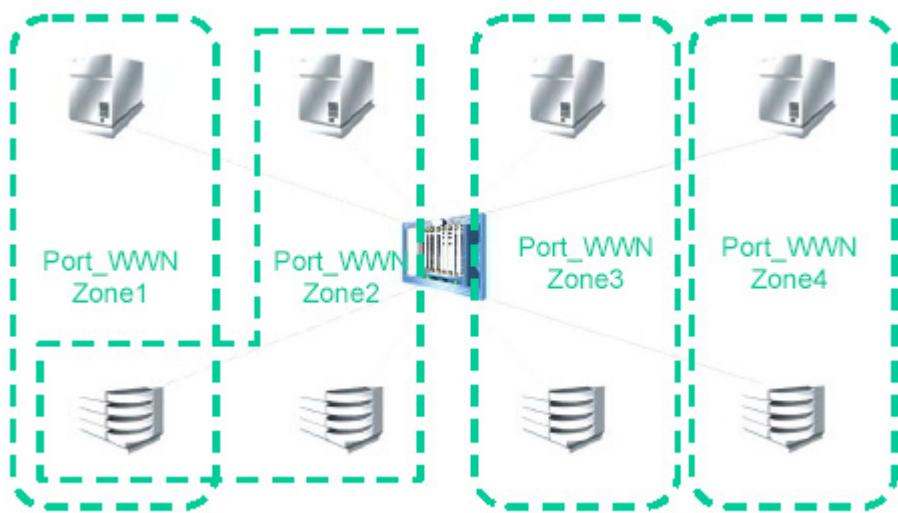


[Figure 1](#) Hardware-enforced non-overlapping zones

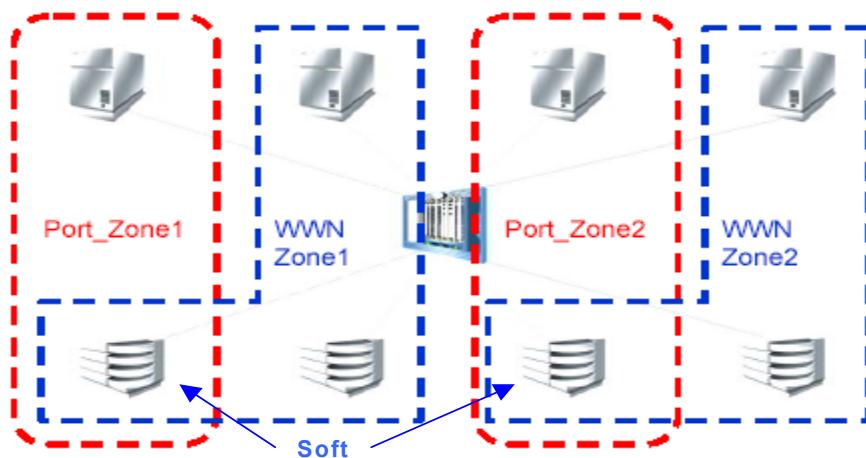
[Figure 2](#) shows the same fabric components zoned in an overlapping fashion.



[Figure 2](#) Hardware-enforced overlapping zones



**Figure 3** Zoning with hardware assist (mixed port and WWN zones)



**Figure 4** Overlapping hardware-enforced zoning with soft porting

In Figure 4, only the ports that are overlapped are software-enforced with hardware assist.

## Rules for configuring zones

Observe the following rules when configuring zones.

- If security is a priority, you should use hard zoning.
- The use of aliases is optional with zoning, and using aliases requires structure when defining zones. However, aliases aid administrators of a zoned fabric to understand the structure and context.
- Evaluate the security requirements of the fabric. If additional security is required, add HP Secure Fabric OS into the fabric.

- If the fabric includes an HP StorageWorks switch and you support a third-party switch product, the third-party switches are able to use only WWN zoning; other types of zoning, including QuickLoop, are not supported.
- QuickLoop

Evaluate whether the fabric will also use QuickLoop Fabric Assist (QLFA) or QuickLoop (QL). If you are running HP Fabric OS v4.x, consider the following before creating and setting up QLFA zones:

- QuickLoop and QuickLoop zones cannot run on switches running Fabric OS v4.x. However, Fabric OS v4.x can still manage (create, remove, update) QuickLoop zones on any non-v4.x switch.
- QuickLoop Fabric Assist. Fabric OS v4.x cannot have a Fabric Assist host directly connected to it. However, targets on a Fabric OS v4.x switch can still be part of a Fabric Assist zone if a Fabric Assist host is connected to a non-v4.x switch.
- Zone changes

Zone changes in a production fabric can cause a disruption of I/O when an RSCN is generated because of the zone change and the HBA is unable to process the RSCN fast enough. Although RSCNs are a normal part of a functioning SAN, the pause in I/O might not be acceptable. For these reasons, you should perform zone changes only when the resulting behavior is predictable and acceptable. Changing HBA drivers can rectify the situation.

- Final verification

After changing or enabling a zone configuration, confirm that the nodes and storage can identify and access one another. Depending on the platform, you might need to reboot one or more nodes in the fabric with the new changes.

The zone configuration is managed on a fabric basis. Zoning can be implemented and administered from any switch in the fabric that has an Advanced Zoning license enabled. When a change in the configuration is saved, enabled, or disabled per the transactional model, it is automatically (by closing the transaction) distributed to all switches in the fabric, preventing a single point of failure for zone information.



**NOTE:** Zoning commands make changes that affect the entire fabric. When executing fabric-level configuration tasks, allow time for the changes to propagate across the fabric before executing any subsequent commands. For a large fabric, you might want to wait several minutes between commands.

## Creating and managing zone aliases

A zone alias is a logical group of ports, WWNs, or AL\_PAs. You can simplify the process of creating zones by first specifying aliases, which eliminates the need for long lists of individual zone member names.

To create an alias:

1. Connect to the switch and log in as admin.
2. Issue the `alcreate` command.
3. Issue the `cfgsave` command to save the change to the defined configuration.

Example:

```
switch:admin> alicreate "array1", "2,32; 2,33; 2,34; 4,4"
switch:admin> alicreate "array2", "21:00:00:20:37:0c:66:23; 4,3"
switch:admin> alicreate "loop1", "4,6[0x02]"
switch:admin> cfgsave
```

To add members to an alias:

1. Connect to the switch and log in as admin.
2. Issue the aliadd command.
3. Issue the cfgsave command to save the change to the defined configuration.

Example:

```
switch:admin> aliadd "array1", "1,2"
switch:admin> aliadd "array2", "21:00:00:20:37:0c:72:51"
switch:admin> aliadd "loop1", "4,6[0x02]"
switch:admin> cfgsave
```

To remove members from an alias:

1. Connect to the switch and log in as admin.
2. Issue the aliremove command.
3. Issue the cfgsave command to save the change to the defined configuration.

Example:

```
switch:admin> aliremove "array1", "1,2"
switch:admin> aliremove "array2", "21:00:00:20:37:0c:72:51"
switch:admin> aliremove "loop1", "4,6[0x02]"
switch:admin> cfgsave
```



**NOTE:** For Fabric OS versions earlier than v4.4.0, when using the aliremove command, the order in which the members appear in the list is critical. For more information on this command, refer to the *HP StorageWorks Fabric OS 4.x command reference guide*.

To delete an alias:

1. Connect to the switch and log in as admin.
2. Issue the alidelete command.
3. Issue the cfgsave command to save the change to the defined configuration.

Example:

```
switch:admin> alidelete "array1"
switch:admin> cfgsave
```

To view an alias in the defined configuration:

1. Connect to the switch and log in as admin.

**2.** Issue the `alishow` command.

The following example shows all zone aliases beginning with arr.

```
switch:admin> alishow "arr*"
alias: array1 21:00:00:20:37:0c:76:8c
alias: array2 21:00:00:20:37:0c:66:23
```

If no parameters are specified, the entire zone database (both the defined and effective configuration) is displayed.

## Creating and maintaining zones

To create a zone:

1. Connect to the switch and log in as admin.
2. Issue the `zonecreate` command.
3. Issue the `cfgsave` command to save the change to the defined configuration.

Example:

```
switch:admin> zonecreate "greenzone", "2,32; 2,33; 2,34; 4,4"
switch:admin> zonecreate "redzone", "21:00:00:20:37:0c:66:23; 4,3"
switch:admin> cfgsave
```

To add devices (members) to a zone:

1. Connect to the switch and log in as admin.
2. Issue the `zoneadd` command.
3. Issue the `cfgsave` command to save the change to the defined configuration.

Example:

```
switch:admin> zoneadd "greenzone", "1,2"
switch:admin> zoneadd "redzone", "21:00:00:20:37:0c:72:51"
switch:admin> zoneadd "bluezone", "4,6[0x02]; 21:00:00:20:37:0c:66:23[0xEF]"
switch:admin> cfgsave
```

To remove devices (members) from a zone:

1. Connect to the switch and log in as admin.
2. Issue the `zoneremove` command.
3. Issue the `cfgsave` command to save the change to the defined configuration.

Example:

```
switch:admin> zoneremove "greenzone", "1,2"
switch:admin> zoneremove "redzone", "21:00:00:20:37:0c:72:51"
switch:admin> zoneremove "bluezone", "4,6[0x02];
21:00:00:20:37:0c:66:23[0xEF]"
switch:admin> cfgsave
```

To delete a zone:

1. Connect to the switch and log in as admin.
2. Issue the `zonedelete` command.
3. Issue the `cfgsave` command to save the change to the defined configuration.

Example:

```
switch:admin> zonedelete "bluezone"
switch:admin> cfgsave
```

To view a zone in the defined configuration:

1. Connect to the switch and log in as admin.
2. Issue the `zoneshow` command.

The following example shows all zones beginning with A, B, or C:

```
switch:admin> zoneshow "[A-C]*"
zone: Blue_zone 1,1; array1; 1,2; array2
zone: Bobs_zone 4,5; 4,6; 4,7; 4,8; 4,9
```

If no parameters are specified, the entire zone database (both the defined and effective configuration) is displayed.

---

## Creating and modifying zoning configurations

You can store a number of zones in a zoning configuration database. The maximum number of items that can be stored in the zoning configuration database depends on the following criteria:

- Number of switches in the fabric.
- Whether or not interoperability mode is enabled.
- Number of bytes per item. The number of bytes required for an item depends on the specifics of the fabric, but cannot exceed 64 bytes per item.

You can use the `cfgsize` command to check both the maximum available size and the currently saved size. See the *HP StorageWorks Fabric OS 4.x command reference guide* for details on the `cfgsize` command. If you believe you are approaching the maximum, you can save a partially completed zoning configuration and use the `cfgsize` command to determine the remaining space.

See “[Managing zoning configurations in a fabric](#)” on page 137 for important considerations for managing zoning in a fabric.

To create a zoning configuration:

1. Connect to the switch and log in as admin.
2. Issue the `cfgcreate` command.
3. Issue the `cfgsave` command to save the change to the defined configuration.

Example:

```
switch:admin> cfgcreate "NEW_cfg", "redzone; bluezone; greenzone"
switch:admin> cfgsave
```

To add zones (members) to a zoning configuration:

1. Connect to the switch and log in as admin.
2. Issue the `cfgadd` command.
3. Issue the `cfgsave` command to save the change to the defined configuration.

Example:

```
switch:admin> cfgadd "newcfg", "bluezone"
switch:admin> cfgsave
```

To remove zones (members) from a zone configuration:

1. Connect to the switch and log in as admin.
2. Issue the `cfgremove` command.
3. Issue the `cfgsave` command to save the change to the defined configuration.

Example:

```
switch:admin> cfgremove "newcfg", "redzone"
switch:admin> cfgsave
```

To delete a zone configuration:

1. Connect to the switch and log in as admin.
2. Issue the `cfgdelete` command.
3. Issue the `cfgsave` command to save the change to the defined configuration.

Example:

```
switch:admin> cfgdelete "testcfg"
switch:admin> cfgsave
```

To clear changes to a configuration:

Use the `cfgtransabort` command. When this command is executed, all changes since the last save operation (performed with the `cfgsave` command) are cleared.

In the following example, assume that the removal of a member from `zone1` was done in error:

```
switch:admin> zoneremove "zone1","3,5"
switch:admin> cfgtransabort
```

To view all zone configuration information:

1. Connect to the switch and log in as admin.
2. Issue the `cfgshow` command with no arguments.

**Example:**

```
switch:admin> cfgshow
Defined configuration:
cfg: USA1 Blue_zone
cfg: USA_cfg Red_zone; Blue_zone
zone: Blue_zone
    1,1; array1; 1,2; array2
zone: Red_zone
    1,0; loop1
alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
alias: loop1 21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df

Effective configuration:
cfg: USA_cfg
zone: Blue_zone
    1,1
    21:00:00:20:37:0c:76:8c
    21:00:00:20:37:0c:71:02
    1,2
    21:00:00:20:37:0c:76:22
    21:00:00:20:37:0c:76:28
zone: Red_zone
    1,0
    21:00:00:20:37:0c:76:85
    21:00:00:20:37:0c:71:df
```

To view selected zone configuration information:

1. Connect to the switch and log in as admin.
2. Issue the **cfgshow** command and specify a pattern.

For example, to display all zone configurations that start with Test:

```
switch:admin> cfgshow "Test*
cfg: Test1 Blue_zone
cfg: Test_cfg Red_zone; Blue_zone
```

To view a configuration in the effective zone database:

1. Connect to the switch and log in as admin.
2. Issue the **cfgactvshow** command.

**Example:**

```
switch:admin> cfgactvshow
Effective configuration:
cfg: NEW_cfg
zone: Blue_zone
    1,1
    21:00:00:20:37:0c:76:8c
    21:00:00:20:37:0c:71:02
    1,2
    21:00:00:20:37:0c:76:22
    21:00:00:20:37:0c:76:28
zone: Red_zone
    1,0
    21:00:00:20:37:0c:76:85
    21:00:00:20:37:0c:71:df
```

# Managing zoning configurations in a fabric

To modify an existing zone configuration, you can add, delete, or remove individual elements to create the desired configuration. After the changes have been made, save the configuration to ensure the configuration is permanently saved in the switch and that the configuration is replicated throughout the fabric.

The switch configuration file can also be uploaded to the host for archiving and it can be downloaded from the host to a switch in the fabric. See “[Backing up a configuration](#)” on page 75 and “[Restoring a configuration](#)” on page 76, or the configupload and configdownload commands in the *HP StorageWorks Fabric OS 4.x command reference guide*.

**Table 21** presents zoning database size limitations for various Fabric OS release versions.

**Table 21** Zoning database limitations

Fabric OS version	Maximum database size (KB)
2.4.0	64
2.5.0	64
2.6.0	96
3.0.0	128
3.1.0	96
3.2.0	256
4.0.0, 4.1.0, 4.2.0	128
4.4.0	256

## Adding a new switch or fabric

When a new switch is added to the fabric, it automatically takes on the zone configuration information from the fabric. Use the cfgactvshow command to verify that the zoning information is the same on each switch in the fabric.

If you are adding a switch that is already configured for zoning, use the cfgclear and cfgsave commands (or use cfgclear and cfgdisable if there is an effective configuration) before connecting the switch to the zoned fabric.

Adding a new fabric that has no zone configuration information to an existing fabric is very similar to adding a new switch. All switches in the new fabric inherit the zoning configuration data. If a zone configuration is in effect, then the same configuration becomes the enabled configuration. The cfgactvshow command displays the same information on all switches in the newly formed fabric.

Before the new fabric can merge successfully, it must pass the following criteria:

- Before merging zones—To facilitate merging, check the following before merging switches or fabrics.
  - Zoning licenses: All switches must have a zoning license enabled.
  - Native operating mode: All switches must be in the native operating mode.
  - HP Secure Fabric OS: If one switch has HP Secure Fabric OS enabled, all switches in the fabric must have HP Secure Fabric OS. Refer to the *HP StorageWorks Secure Fabric OS user guide* for more information.

- Merging and segmentation—The fabric is checked for segmentation during power-up or when a switch is disabled or enabled, or when a new switch is added.

Two databases are used with zoning. The first database is the zone configuration database. This is the data displayed as the defined configuration in the `cfgShow` command. It is stored in nonvolatile memory by the `cfgSave` command. This database is a replicated database, which means that all switches in the fabric has a copy of this database. When a change is made to the defined configuration, the switch where the changes were made must close its transaction for the change to get propagated throughout the fabric.

- Merging rules—Observe the following rules when merging zones:

Local and adjacent configurations	If the local and adjacent zone database configurations are the same, they remain unchanged after the merge.
Effective configurations	If there is an effective configuration between two switches, the zone configuration in effect match.
Zone object naming	If a zoning object has the same name in both the local and adjacent defined configurations, the object types and member lists must match. When comparing member lists, the content and order of the members are important.
Objects in adjacent configurations	If a zoning object appears in an adjacent defined configuration, but not in the local defined configuration, the zoning object is added to the local defined configuration. The modified zone database must fit in the nonvolatile memory area allotted for the zone database.
Local configuration modification	If a local defined configuration is modified because of a merge, the new zone database is propagated to other the switches within the merge request.

- Merging two fabrics—Both fabrics have identical zones and configurations enabled. The two fabrics join to make one larger fabric with the same zone configuration across the newly created fabric.

If the two fabrics have different zoning configurations, they are merged. If the two fabrics cannot join, the ISL between the switches are segmented.

- Merge conflicts—When a merge conflict is present, a merge does not take place and the ISL segments. Use the `switchshow` command to obtain additional information about possible merge conflicts, because many non-zone related configuration parameters can cause conflicts

If the fabrics have different zone configuration data, the system attempts to merge the two sets of zone configuration data. If the zones cannot merge, the ISL segmenteds.

A merge is not possible if any of the following conditions exist:

Configuration mismatch	Zoning is enabled in both fabrics and the zone configurations that are enabled are different in each fabric.
Type mismatch	The name of a zone object in one fabric is used for a different type of zone object in the other fabric.
Content mismatch	The definition of a zone object in one fabric is different from the definition of zone object with the same name in the other fabric.



**NOTE:** If the zoneset members on two switches are not listed in the same order, the configuration is considered a mismatch, which causes the switches to be segmented from the fabric. For example: cfg1 = z1; z2 is different from cfg1 = z2; z1, even though members of the configuration are the same. If zoneset members on two switches have the same names defined in the configuration, make sure zoneset members are listed in the same order.

## Splitting a fabric

If the connections between two fabrics are no longer available, the fabric will segment into two separate fabrics. Each new fabric retains the same zone configuration.

If the connections between two fabrics are replaced and no changes have been made to the zone configuration in either of the two fabrics, then the two fabrics merge back into one single fabric. If any changes that cause a conflict have been made to either zone configuration, then the fabrics might segment.

## Using zoning to administer security

Zones provide controlled access to fabric segments and establish barriers between operating environments. They isolate systems with different uses, protecting individual systems in a heterogeneous environment; for example, when zoning is in secure mode, no merge operations occur.

HP Advanced Zoning is configured on the primary Fabric Configuration Server (FCS). The primary FCS switch makes zoning changes and other security-related changes. The primary FCS switch also distributes zoning to all other switches in the secure fabric. All existing interfaces can be used to administer zoning (depending on the policies; refer to the *HP StorageWorks Secure Fabric OS user guide* for information about security policies).

You must perform zone management operations from the primary FCS switch using a zone management interface, such as telnet or Advanced Web Tools. You can alter a zoning database, provided you are connected to the primary FCS switch.

When two secure fabrics join, the traditional zoning merge does not occur. Instead, a zoning database is downloaded from the primary FCS switch of the merged secure fabric. When E\_Ports are active between two switches, the name of the FCS server and a zoning policy set version identifier are exchanged between the switches. If the views of the two secure fabrics are the same, the fabric's primary FCS server downloads the zoning database and security policy sets to each switch in the fabric. If there is a view conflict, the E\_Ports are segmented due to incompatible security data.

As part of zoning architecture, you must determine which of the two basic zoning architectures (hard or soft) works best for your fabric. With time and planning, the basic hard zone configuration works for most sites. If a site has additional security needs, use the additional layer of Secure Fabric OS, apart from the standard zoning architecture.



**NOTE:** Secure Fabric OS requires the activation of a HP security license and an Advanced Zoning license.

# Resolving zone conflicts

Zone conflicts can be resolved by saving a configuration file with the configupload command, examining the zoning information in the file, and performing a cut and paste operation so that the configuration information matches in the fabrics being merged.

After examining the configuration file, you can choose to resolve zone conflicts by using the cfgclear command followed by the cfgdisable command on the incorrectly configured segmented fabric, followed by a portdisable or portenable command on one of the ISL ports that connects the fabrics. This causes a merge, making the fabric consistent with the correct configuration. [Table 22](#) lists considerations for zoning architecture.



**CAUTION:** Be careful using the cfgclear command; it deletes the defined configuration.

**Table 22** Considerations for zoning architecture

Item	Description
Type of zoning: hard or soft (session-based)	If security is a priority, HP recommends hard zoning.
Use of aliases	The use of aliases is optional with zoning. Using aliases requires structure when defining zones. Aliases aid administrators of zoned fabric in understanding the structure and context.
Security requirements	Evaluate the security requirements of the fabric. If additional security is required, add HP Secure Fabric OS into the fabric.
Interoperability fabric	If the fabric includes a third-party switch product, only WWN zoning is supported. Other types of zoning, including QuickLoop, are not supported.
QLFA zones	Evaluate whether the fabric will have QuickLoop Fabric Assist (QLFA) or QuickLoop (QL) in it, and consider the following items before creating and setting up QLFA zones:  QuickLoop Zoning—QuickLoop/QuickLoop zones cannot run on Fabric OS v4.1.0 or later. However, Fabric OS can manage (create, remove, update) QL zones.  QuickLoop Fabric Assist—A switch running Fabric OS v4.1.0 or later cannot have a Fabric Assist host directly connected to it. However, such a switch can be part of a Fabric Assist zone if a Fabric Assist host is connected to a compatible switch in the fabric.
Testing	Testing a (new) zone configuration—Before implementing a zone, run the Zone Analyzer from Advanced Web Tools to isolate any possible problems. This is especially useful as fabrics increase in size.

**Table 22** Considerations for zoning architecture (continued)

Item	Description
Effect of changes in a production fabric	Zone changes in a production fabric can result in a disruption of I/O under conditions where an RSCN is issued as a result of a zone change and the HBA is unable to process the RSCN fast enough. Though RSCNs are a normal part of a functioning SAN, the pause in I/O may not be acceptable. For these reasons, HP recommends that you perform zone changes only when the resulting behavior is predictable and acceptable. Changing HBA drivers can correct the situation.
Confirming operation	After changing or enabling a zone configuration, confirm that the nodes and storage are able to identify and access one another. Depending on the platform, you may need to reboot one or more nodes in the fabric with the new changes.



# 10 Administering advanced performance monitoring

This chapter contains procedures for the HP Advanced Performance Monitoring licensed feature and contains the following sections:

- [Displaying and clearing the CRC error count](#), page 145
- [Monitoring end-to-end performance](#), page 145
- [Monitoring filter-based performance](#), page 149
- [Monitoring ISL performance](#), page 152
- [Monitoring trunks](#), page 152
- [Displaying monitor counters](#), page 153
- [Clearing monitor counters](#), page 155
- [Saving and restoring monitor configurations](#), page 156
- [Collecting performance data](#), page 156

Based on HP Frame Filtering technology and a unique performance counter engine, Advanced Performance Monitoring is a comprehensive tool for monitoring the performance of networked storage resources. It supports direct-attach, loop, and switched fabric Fibre Channel SAN topologies by:

- Monitoring transaction performance from source to destination
- Reporting cyclic redundancy check (CRC) error measurement statistics
- Measuring HP Inter-Switch Link (ISL) Trunking performance and resource usage

Further features are provided through HP Advanced Web Tools:

- Measuring device performance by port, arbitrated loop physical address (AL\_PA), and logical unit number (LUN)
- Comparing IP versus SCSI traffic on each port
- Providing a library of predefined graphs



**NOTE:** The SAN Switch 4/32 running Fabric OS v4.4.0 does not display AL\_PA measurements for end-to-end monitors. It provides port CRC reports through Advanced Web Tools.

[Table 23](#) lists commands associated with Advanced Performance Monitoring. For detailed information on these commands, refer to the *HP StorageWorks Fabric OS 4.x command reference guide*.

**Table 23** Advanced performance monitoring commands

Command	Description
perfaddeemonitor	Add an end-to-end (EE) monitor to a port.
perfaddipmonitor	Add an IP monitor to a port.
perfaddrademonitor	Add a SCSI Read monitor to a port.
perfaddrwmonitor	Add a SCSI Read and Write monitor to a port.
perfaddscsimonitor	Add a SCSI traffic frame monitor to a port.
perfaddusermonitor	Add a user-defined monitor to a port.
perfaddwritemonitor	Add a SCSI Write monitor to a port.
perfcfgclear	Clear the performance monitoring settings from nonvolatile memory.
perfcfgrestore	Restore performance monitoring settings from nonvolatile memory.
perfcfgsave	Save the current performance monitoring settings to nonvolatile memory.
perfclralpacrc	Clear an AL_PA device CRC count by the port and AL_PA.
perfdeleemonitor	Delete an end-to-end monitor on port.
perfdefiltermonitor	Delete a filter-based monitor.
perfmonitordclear	Clear statistics counters of end-to-end, filter-based, and ISL monitors on a port.
perfmonitorshow	Display end-to-end, filter-based, and ISL monitors on a port.
perfsetportemask	Set overall mask for end-to-end monitors.
perfshowalpacrc	Display the AL_PA CRC count by port or by AL_PA.
perfshowportemask	Display the current end-to-end mask of a port.



**NOTE:** The command examples use the slot/port syntax required by the Core Switch 2/64 and the SAN Director 2/128. For HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32, use only the port number where needed in the commands.

# Displaying and clearing the CRC error count

You can use the `perfshowalpacrc` command to display the CRC error count for all AL\_PA devices or a single AL\_PA on a specific active L\_Port.



**NOTE:** The SAN Switch 4/32 running Fabric OS v4.4.0 provides port CRC reports through Advanced Web Tools.

To display the CRC error count for all AL\_PA devices on a port:

```
switch:admin> perfshowalpacrc 1/1
AL_PA      CRC count
-----
0xd9          0
```

To display the CRC error count for a single AL\_PA device on a port:

```
switch:admin> perfshowalpacrc 1/1, 0xd9
The CRC count at ALPA 0xd9 on port 1 is 0x0000000000.
```

To clear the CRC error count:

```
switch:admin> perfclralpacrc 1/1, 0xd9
CRC error count at AL_PA 0xd9 on port 1 is cleared.
switch:admin> perfclralpacrc 1/1
No AL_PA value is specified. This will clear all AL_PA CRC
counts on port 1. Do you want to continue? (yes, y, no, n): [no] y
Please wait ...
All alpa CRC counts are cleared on port 1.
```

In v3.1.0, v4.1.0, and later, the `portstatsclear` command clears AL\_PA-based CRC error counters for all the ports in the same group.

# Monitoring end-to-end performance

End-to-end performance monitoring counts the number of words and CRC errors in Fibre Channel frames for a specified Source ID (SID) and Destination ID (DID) pair. An end-to-end performance monitor includes these counts:

- RX\_COUNT (words in frames received at the port)
- TX\_COUNT (words in frames transmitted from the port)
- CRC\_COUNT (frames with CRC errors received at or transmitted from the port)

To enable end-to-end performance monitoring, you must configure an end-to-end monitor on a port, specifying the SID-DID pair (in hexadecimal). The monitor counts only those frames with matching SID and DID.

Each SID or DID has three fields, listed in the following order:

- Domain ID (DD)
- Area ID (AA)
- AL\_PA (PP)

For example, the SID 0x118a0f denotes DD 0x11, AA 0x8a, and AL\_PA 0x0f.

You can monitor end-to-end performance using the `perfmonitorshow` command, as described in “[Displaying monitor counters](#)” on page 153. You can clear end-to-end counters using the `perfmonitorclear` command, as described in “[Clearing monitor counters](#)” on page 155.



**NOTE:** For end-to-end monitors, CRC counters are not displayed on the SAN Switch 4/32.

## Adding end-to-end monitors

An end-to-end monitor counts the following items for a port: number of words received, number of words transmitted, and number of CRC errors detected in frames.

The HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128 allow up to eight end-to-end monitors.

The SAN Switch 4/32 allows up to 256 end-to-end monitors shared by all ports. (The number of interswitch links configured on the switch affects the amount of resources available for end-to-end monitors.)

End-to-end monitors cannot be added to interswitch links.

The monitor count is qualified using either of following conditions:

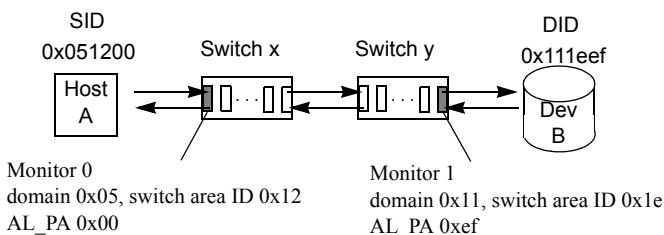
- For frames received at the port with the end-to-end monitor installed, the frame SID is the same as `SourceID` and the frame DID is the same as `DestID`. The `RX_COUNT` and `CRC_COUNT` are updated accordingly.
- For frames transmitted from the port with the end-to-end monitor installed, the frame DID is the same as `SourceID` and the frame SID is the same as `DestID`. The `TX_COUNT` and `CRC_COUNT` are updated accordingly.



**NOTE:** How the area ID for a port relates to the port number depends upon the PID format used by the fabric. See “[Configuring the PID format](#)” on page 203 for more information.

**Figure 5** shows two devices:

- Host A is connected to domain 5 (0x05), switch area ID 18 (0x12), AL\_PA 0x00 on Switch X.
- Dev B is a storage device connected to domain 17 (0x11), switch area ID 30 (0x1e), AL\_PA 0xef on Switch Y.



**Figure 5** Setting end-to-end monitors on a port



**NOTE:** End-to-end performance monitoring looks at traffic on the receiving port respective to the SID only. In [Figure 5](#), if you add a monitor to slot 2, port 2 on Switch X, specifying Dev B as the SID and Host A as the DID, no counters (except CRC) are incremented.

To monitor the traffic from Host A to Dev B:

```
switch:admin> perfaddeemonitor 2/2, "0x051200" "0x111eef"  
End-to-End monitor number 0 added.
```

Add Monitor 0 to slot 2, port 2 on Switch X, specifying 0x051200 as the SID and 0x111eef as the DID, as shown in the following example:

Monitor 0 counts the frames that have an SID of 0x051200 and a DID of 0x111eef. For monitor 0, RX\_COUNT is the number of words from Host A to Dev B, TX\_COUNT is the number of words from Dev B to Host A, and CRC\_COUNT is the number of frames in both directions with CRC errors.

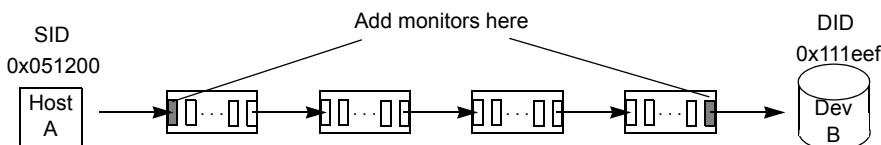
To monitor the traffic from Dev B to Host A:

```
switch:admin> perfaddeemonitor 2/14, "0x111eef" "0x051200"  
End-to-End monitor number 1 added.
```

Add Monitor 1 to slot 2, port 14 on Switch y, specifying 0x111eef as the SID and 0x051200 as the DID, as shown in the following example.

Monitor 1 counts the frames that have an SID of 0x111eef and a DID of 0x051200. For Monitor 1, RX\_COUNT is the number of words from Dev B to Host A, TX\_COUNT is the number of words from Host A to Dev B, and CRC\_COUNT is the number of frames in both directions with CRC errors.

**Figure 6** shows several switches and the correct ports on which to add performance monitors for a specified SID-DID pair.



**Figure 6** Proper placement of end-to-end performance monitors

## Setting a mask for end-to-end monitors

End-to-end monitors count the number of words in Fibre Channel frames that match a specific SID/DID pair. If you want to match only part of the SID or DID, you can set a mask on the port to compare only certain parts of the SID or DID. By default, the frame must match the entire SID and DID to trigger the monitor. By setting a mask, you can choose to have the frame match only one or two of the three fields (Domain ID, Area ID, and AL\_PA) to trigger the monitor.



**NOTE:** Only one mask per port can be set. When you set a mask, all existing end-to-end monitors are deleted.

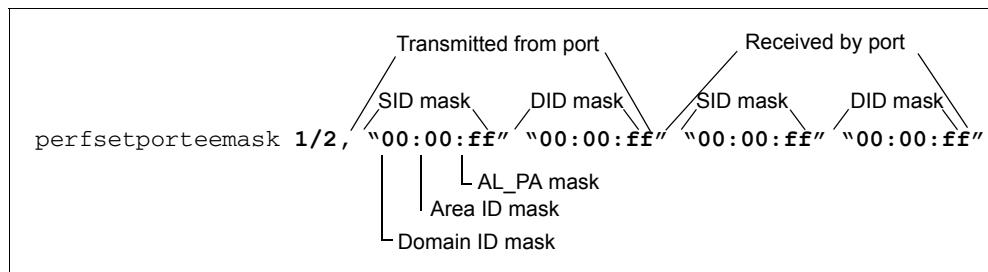
You can specify a mask using the `perfsetportemask` command in the form `dd:aa:pp`, where `dd` is the Domain ID mask, `aa` is the Area ID mask, and `pp` is the AL\_PA mask. The values for `dd`, `aa`, and `pp` are either `ff` (the field must match) or `00` (the field is ignored). The default EE mask value is `ff:ff:ff`. The command sets the mask for all end-to-end monitors of a port. If any end-to-end monitors are programmed on a port when the `perfsetportemask` command is issued, a message appears similar to that in the following example:

```
switch:admin> perfsetportemask 1/2, "00:00:ff"
EE monitors are currently programmed on this port. Changing EE mask
for this port will cause ALL EE monitors on this port to be deleted.
Do you want to continue? (yes, y, no, n): [no] y

EE mask on port <port-number> is set and EE monitors were deleted
```

The `perfsetportemask` command sets a mask for the Domain ID, Area ID, and AL\_PA of the SIDs and DIDs for frames transmitted from and received by the port.

Figure 7 shows the mask positions in the command. A mask (`ff`) is set on slot 1, port 2 to compare the AL\_PA fields on the SID and DID in all frames (transmitted and received) on port 2. The frame SID and DID must match only the AL\_PA portion of the specified SID-DID pair. Each port can have only one EE mask. The mask is applied to all end-to-end monitors on the port. Individual masks for each monitor on the port cannot be specified.



**Figure 7** Mask positions for end-to-end monitors

To display the current end-to-end mask of a port, use the `perfshowportemask` command. The end-to-end mask has 12 fields, each with a value of on or off.

To set and display an end-to-end mask:

```
switch:admin> perfsetportemask 1/11, "00:00:ff" "00:00:ff" "00:00:ff"  
"00:00:ff"  
The EE mask on port 11 is set and EE counters are reset.  
switch:admin> perfshowportemask 1/11  
The EE mask on port 11 is set by application TELNET  
TxSID Domain: off  
TxSID Area: off  
TxSID AL_PA: on  
TxDID Domain: off  
TxDID Area: off  
TxDID AL_PA: on  
RxSID Domain: off  
RxSID Area: off  
RxSID AL_PA: on  
RxDID Domain: off  
RxDID Area: off  
RxDID AL_PA: on
```

The end-to-end mask is set on slot 1, port 11.

To display a monitor, use the `perfmonitorshow` command, as described in “[Displaying monitor counters](#)” on page 153.

## Deleting end-to-end monitors

Use the `perfdeleemonitor` command to delete end-to-end monitors. You can delete all monitors or specific monitors. The following example deletes the end-to-end monitor number 0 on slot 1, port 2:

```
switch:admin> perfdeleemonitor 1/2, 0  
End-to-End monitor number 0 deleted
```

---

## Monitoring filter-based performance

Filter-based performance monitoring counts the number of times a frame with a particular pattern is transmitted by a port. Filter-based monitoring is achieved by configuring a filter for a particular purpose. The filter can be a standard filter (for example, a SCSI Read command filter that counts the number of SCSI Read commands that have been transmitted by the port) or a user-defined filter customized for your particular use.

For HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128, the maximum number of filters is eight per port, in any combination of standard filters and user-defined filters.

For the SAN Switch 4/32, the maximum number of filters is 12 per port, in any combination of standard filters and user-defined filters.

The actual number of filters that can be configured on a port depends on the complexity of the filters. For trunked ports, the filter is configured on the trunk master.

You can monitor filter-based performance using the `perfmonitorshow` command, as described in “[Displaying monitor counters](#)” on page 153. You can clear filter-based counters using the `perfmonitorclear` command, as described in “[Clearing monitor counters](#)” on page 155.

## Adding standard filter-based monitors

Table 24 lists the commands for adding standard filter-based monitors to a port.

**Table 24** Commands to add filter-based monitors

Telnet command	Description
perfaddreadmonitor	Count the number of SCSI Read commands.
perfaddwritemonitor	Count the number of SCSI Write commands.
perfaddrwmonitor	Count the number of SCSI Read and Write commands.
perfaddscsimonitor	Count the number of SCSI traffic frames.
perfaddipmonitor	Count the number of IP traffic frames.

The following example adds filter-based monitors to slot 1, port 2 and displays the results:

```
switch:admin> perfaddreadmonitor 1/2
SCSI Read filter monitor #0 added
switch:admin> perfaddwritemonitor 1/2
SCSI Write filter monitor #1 added
switch:admin> perfaddrwmonitor 1/2
SCSI Read/Write filter monitor #2 added
switch:admin> perfaddscsimonitor 1/2
SCSI traffic frame monitor #3 added
switch:admin> perfaddipmonitor 1/2
IP traffic frame monitor #4 added
switch:admin> perfmonitorshow --class FLT 1/2
There are 5 filter-based monitors defined on port 2.
```

KEY	ALIAS	OWNER_APP	OWNER_IP_ADDR	FRAME_COUNT
0	SCSI Read	TELNET	N/A	0x0000000000000000
1	SCSI Write	TELNET	N/A	0x0000000000000000
2	SCSI R/W	TELNET	N/A	0x0000000000000000
3	SCSI Frame	TELNET	N/A	0x0000000000000000
4	IP Frame	TELNET	N/A	0x0000000000000000

## Adding custom filter-based monitors

In addition to the standard filters—read, write, read/write, SCSI frame and IP frame—you can create custom filters to gather statistics that fit your needs.

To define a custom filter, use the `perfaddusermonitor` command. With this command, you must specify a series of *offsets*, *masks*, and *values*. For all transmitted frames, the switch performs these tasks:

- Locates the byte found in the frame at the specified *offset*.
- Applies the *mask* to the byte found in the frame.
- Compares the value with the given *values* in the `perfaddusermonitor` command.
- Increments the filter counter if a match is found.

The following number of offsets can be specified:

- For the HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128 (Fabric OS v4.0.0 or later), up to two different offsets per port.
- For the SAN Switch 2/8-EL and SAN Switch 2/16 (Fabric OS v3.0.0 or later), up to three different offsets per port.
- For the SAN Switch 4/32 (Fabric OS v4.4.0 or later), up to 15 different offsets per port (14 offsets when FMS is enabled).

You can specify up to four values to compare against each offset. If more than one offset is required to properly define a filter, the bytes found at each offset must match one of the given values for the filter to increment its counter. If one or more of the given offsets does not match any of the given values, the counter does not increment. The value of the offset must be between 0 and 63, in decimal format. Byte 0 indicates the first byte of the Start of Frame (SOF), byte 4 is the first byte of the frame header, and byte 28 is the first byte of the payload. Thus only the SOF, frame header, and first 36 bytes of payload can be selected as part of a filter definition. Offset 0 is a special case, which can be used to monitor the first 4 bytes of the frame (SOF). When the offset is set to 0, the values 0–7 that are checked against that offset are predefined as shown in [Table 25](#).

**Table 25** Predefined values at offset 0

Value	SOF
0	SOFF
1	SOFc1
2	SOFi1
3	SOFn1
4	SOFi2
5	SOFn2
6	SOFi3
7	SOFn3

If the switch does not have enough resources to create a given filter, you may have to delete other filters to free resources.

To add filter-based monitors:

```
switch:admin> perfaddusermonitor 4/2, "12, 0xff, 0x05, 0x08; 9, 0xff, 0x02" "FCP/IP"
User monitor #5 added
switch:admin> perfaddusermonitor 1/2, "0, 0xff, 6"
User Monitor #6 added
```

In the previous example, two filter-based monitors are added. The first monitor (#5) counts all FCP and IP frames transmitted from domain 0x02 for slot 4, port 2. The FCP and IP protocols are selected by monitoring offset 12, mask 0xff and matching values of 0x05 or 0x08. Domain 2 is selected by monitoring offset 9, mask 0xff, and matching a value of 0x02. The monitor counter is incremented for all outgoing frames from port 2 where byte 9 is 0x02 and byte 12 is 0x05 or 0x08.

The second monitor (#6) is for SOFi3 on slot 1, port 2.

## Deleting filter-based monitors

To delete a filter-based monitor:

1. List the valid monitor numbers using the `perfshowfiltermonitor` command.
2. Use the `perfdelfiltermonitor` command to delete a specific monitor. If you do not specify which monitor number to delete, you are asked if you want to delete all entries.

The following example displays the monitors on slot 1, port 4 using the `perfshowfiltermonitor` command (the monitor numbers are listed in the KEY column) and deletes monitor number 1 on slot 1, port 4 using the `perfdelfiltermonitor` command:

```
switch:admin> perfshowfiltermonitor 1/4
There are 4 filter-based monitors defined on port 4.
KEY      ALIAS      OWNER_APP      OWNER_IP_ADDR      FRAME_COUNT
-----
0  SCSI  Read    TELNET          N/A            0x00000000000002208
1  SCSI  Write   TELNET          N/A            0x0000000000000464a
2  SCSI  R/W     TELNET          N/A            0x0000000000000fd8c
3  SCSI  Frame   WEB_TOOLS     192.168.169.40  0x000000000002c2229
switch:admin> perfdelfiltermonitor 1/4, 1
The specified filter-based monitor is deleted.
```

---

## Monitoring ISL performance

ISL monitoring is set up on E\_Ports automatically in release v4.4.0 and later.

An ISL monitor measures traffic to all reachable destination domains for an ISL, showing which destination domain is consuming the most traffic. If there are more than 16 domains, the monitor samples traffic and extrapolates the measurement.

You can monitor ISL performance using the `perfmonitorshow` command, as described in ["Displaying monitor counters"](#) on page 153. You can clear ISL counters using the `perfmonitorclear` command, as described in ["Clearing monitor counters"](#) on page 155.

---

## Monitoring trunks

For trunked ISLs on Fabric OS v4.x switches, monitoring is set only on the master ISL, which communicates with the associated slave ISLs. For Fabric OS v3.x switches, monitoring can be set on slave ISLs.

End-to-end monitors are not supported for ISLs.

The HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128 support eight filter-based monitors for trunks.

The SAN Switch 4/32 supports 12 filter-based monitors for trunks.

# Displaying monitor counters

Use the `perfmonitorshow` command to display the monitors on a specified port. For end-to-end counters, you can display either the cumulative count of the traffic detected by the monitors or a snapshot of the traffic at specified intervals.



**NOTE:** The SAN Switch 4/32 output does not include CRC counts.

The command format is:

```
perfmonitorshow --class monitor_class [slotnumber/]portnumber [interval]
```

where:

*monitor\_class* Specifies the monitor class, which can be `EE` (end-to-end), `FLT` (filter-based), or `ISL` (inter-switch link). The `--class monitor_class` operand is required.

*slotnumber* Specifies the slot number for a Core Switch 2/64 and SAN Director 2/128. For all other switches, this operand is not required. The slot number must be followed by a slash ( / ) and the port number, so that each port is represented by both slot number (1 through 4 or 7 through 10) and port number (0 through 15).

The Core Switch 2/64 and SAN Director 2/128 each has a total of 10 slots. Slot numbers 5 and 6 are control processor cards; slots 1 through 4 and 7 through 10 are port cards. On each port card, there are 16 ports, counted from the bottom, numbered 0 to 15.

*portnumber* Specifies a port number. Valid values for port number vary, depending on the switch type. This operand is required.

*interval* Specifies an interval in seconds. The interval must be greater than or equal to 5 seconds. For end-to-end monitoring, the Tx and Rx counts are measured in bytes. This operand is optional.

To display an end-to-end monitor on a port at an interval of every 6 seconds:

```
switch:admin> perfMonitorShow --class EE 4/5 6
perfmonitorshow 53, 6: Tx/Rx are # of bytes and crc is # of crc errors
      0           1           2           3           4
----- -----
crc  Tx   Rx   crc  Tx   Rx   crc  Tx   Rx   crc  Tx   Rx   crc  Tx   Rx
===== ===== ===== ===== ===== ===== ===== ===== ===== ===== ===== ===== =====
0    0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
0    53m  4.9m 0
0    53m  4.4m 0
0    53m  4.8m 0
0    53m  4.6m 0
0    53m  5.0m 0
0    53m  4.8m 0
0    53m  4.5m 0
0    52m  4.5m 0
0    52m  5.0m 0
0    52m  4.5m 0
0    52m  4.6m 0
```

To display EE monitors on a port:

KEY	SID	DID	OWNER_APP	OWNER_IP_ADDR	TX_COUNT	RX_COUNT	CRC_COUNT
0	0x58e0f	0x1182ef	TELNET	N/A	0x0000000000000000	0x0000000000000000	0x0000000000000000
0	0x21300	0x21dda	TELNET	N/A	0x00000004d0ba9915	0x000000067229e65	0x0000000000000000
1	0x21300	0x21ddc	TELNET	N/A	0x00000004d0baa754	0x000000067229e65	0x0000000000000000
2	0x21300	0x21de0	TELNET	N/A	0x00000004d0bab3a5	0x000000067229e87	0x0000000000000000
3	0x21300	0x21de1	TELNET	N/A	0x00000004d0bac1e4	0x000000067229e87	0x0000000000000000
4	0x21300	0x21de2	TELNET	N/A	0x00000004d0bad086	0x000000067229e87	0x0000000000000000
5	0x11000	0x21fd6	WEB_TOOLS	192.168.169.40	0x00000004d0bade54	0x000000067229e87	0x0000000000000000
6	0x11000	0x21fe0	WEB_TOOLS	192.168.169.40	0x00000004d0baed41	0x000000067229e98	0x0000000000000000

To display a filter-based monitor on a port at an interval of every 6 seconds:

switch:admin> <b>perfMonitorShow --class FLT 2/5 6</b>							
perfmonitorshow 21, 6							
0	1	2	3	4	5	6	
#Frames	#Frames	#Frames	#Frames	#Frames	#Frames	#Frames	
0	0	0	0	0	0	0	
26k	187	681	682	682	494	187	
26k	177	711	710	710	534	176	
26k	184	734	734	734	550	184	
26k	182	649	649	649	467	182	
26k	188	754	755	755	567	184	
26k	183	716	716	717	534	183	
26k	167	657	656	655	488	167	
26k	179	749	749	749	570	179	
26k	164	752	752	752	588	164	
26k	190	700	700	700	510	190	
26k	181	701	701	701	520	181	
26k	200	750	750	751	550	201	
26k	180	692	692	691	512	179	
26k	179	696	696	696	517	179	
26k	187	720	720	720	533	187	
26k	200	722	722	722	522	200	
26k	204	717	717	717	513	204	

To display a filter monitor information on a port:

switch:admin> <b>perfMonitorShow --class FLT 2/5</b>				
There are 7 filter-based monitors defined on port 21.				
KEY	ALIAS	OWNER_APP	OWNER_IP_ADDR	FRAME_COUNT
0	SCSI_Frame	TELNET	N/A	0x00000000002c2229
1	SCSI_WR	TELNET	N/A	0x000000000000464a
2	SCSI_RW	TELNET	N/A	0x000000000000fd8c
3	SCSI_RW	WEB_TOOLS	192.168.169.40	0x0000000000007ba3
4	SCSI_RW	WEB_TOOLS	192.168.169.190	0x0000000000004f0e
5	SCSI_RD	WEB_TOOLS	192.168.169.40	0x0000000000002208
6	SCSI_WR	WEB_TOOLS	192.168.169.40	0x00000000000033a

To display an ISL monitor information on a port:

```
switch:admin> perfMonitorShow --class ISL 1/1
Total transmit count for this ISL: 1462326
Number of destination domains monitored: 3
Number of ports in this ISL: 2
Domain 97:           110379          Domain 98:           13965
Domain 99:           1337982
```

## Clearing monitor counters

Before you clear statistics counters, verify the valid monitor numbers on a specific port using the **perfmonitorshow** command, to make sure the correct monitor counters are cleared. To clear statistics counters for all or a specified monitor, use the **perfmonitorclear** command. After the command has been executed, the telnet shell confirms that the counters on the monitor have been cleared.

The command format is:

```
perfmonitorclear --class monitor_class [slotnumber/]portnumber [monitorId]
```

where:

*monitor\_class* Specifies the monitor class, which can be EE (end-to-end), FLT (filter-based), or ISL (inter-switch link). The **--class monitor\_class** operand is required.

*slotnumber* Specifies the slot number for a Core Switch 2/64 or SAN Director 2/128. For all other switches, this operand is not required. The slot number must be followed by a slash ( / ) and the port number, so that each port is represented by both slot number (1 through 4 or 7 through 10) and port number (0 through 15).

The Core Switch 2/64 or SAN Director 2/128 each has a total of 10 slots. Slot numbers 5 and 6 are control processor cards; slots 1 through 4 and 7 through 10 are port cards. On each port card, there are 16 ports, counted from the bottom, numbered 0 to 15.

*portnumber* Specifies a port number. Valid values for port number vary, depending on the switch type. This operand is required.

*monitorId* Specifies the monitor number to clear. Monitor numbers are defined when you create the monitor on a port. This operand is optional. If not specified, all monitor counters on the port are cleared. This operand does not apply to ISL monitors.



**NOTE:** In Fabric OS v3.1.0 and v4.1.0 (or later) the **portstatsclear** command clears AL\_PA-based CRC error counters for all the ports in the same group.

To clear statistics counters for an end-to-end monitor:

```
switch:admin> perfMonitorClear --class EE 1/2 5
End-to-End monitor number 5 counters are cleared

switch:admin> perfMonitorClear --class EE 1/2
This will clear ALL EE monitors' counters on port 2, continue?
(yes, y, no, n): [no] y
```

To clear statistics counters for a filter-based monitor:

```
switch:admin> perfMonitorClear --class FLT 1/2 4
Filter-based monitor number 4 counters are cleared

switch:admin> perfMonitorClear --class FLT 1/2
This will clear ALL filter-based monitors' counters on port 2, continue?
(yes, y, no, n): [no] y
```

To clear statistics counters for an ISL monitor:

```
switch:admin> perfMonitorClear --class ISL 1
This will clear ISL monitor on port 1, continue? (yes, y, no, n): [no] y
```

---

## Saving and restoring monitor configurations

To save the current end-to-end and filter monitor configuration settings into nonvolatile memory, use the **perfcfgsave** command; for example:

```
switch:admin> perfcfgsave
This will overwrite previously saved Performance Monitoring settings in
FLASH ROM. Do you want to continue? (yes, y, no, n): [no] y
Please wait... Committing configuration...done.
Performance monitoring configuration saved in FLASH ROM.
```

To restore a saved monitor configuration, use the **perfcfgrestore** command; for example, to restore the original performance monitor configuration after making several changes:

```
switch:admin> perfcfgrestore
This will overwrite current Performance Monitoring settings in RAM. Do
you want to continue? (yes, y, no, n): [no] y
Please wait... Performance monitoring configuration restored from FLASH
ROM.
```

To clear the previously saved performance monitoring configuration settings from nonvolatile memory, use the **perfcfgclear** command; for example:

```
switch:admin> perfcfgclear
This will clear Performance Monitoring settings in FLASH ROM. The RAM
settings won't change. Do you want to continue? (yes, y, no, n): [no] y
Please wait... Committing configuration...done.
Performance Monitoring configuration cleared from FLASH.
```

---

## Collecting performance data

Data collected through Advanced Performance Monitoring is deleted when the switch is rebooted. Using the HP Fabric Manager software application version 4.4.0 (or later), you can store performance data persistently. For details on this feature, refer to the *HP StorageWorks Fabric Manager 4.4.x user guide*.

# 11 Configuring the distributed management server

This chapter contains the following sections:

- [Enabling and disabling the platform services](#), page 157
- [Controlling access](#), page 158
- [Configuring the server database](#), page 161
- [Controlling topology discovery](#), page 162

The HP Fabric OS Distributed Management Server allows a SAN management application to retrieve information and administer interconnected switches, servers, and storage devices. The management server assists in the auto-discovery of switch-based fabrics and their associated topologies.

A client of the management server can find basic information about the switches in the fabric and use this information to construct topology relationships. The management server also allows you to obtain certain switch attributes and, in some cases, modify them. For example, logical names identifying switches can be registered with the management server.

## Enabling and disabling the platform services

The management server is located at the Fibre Channel well-known address FFFFFAh. All management services except platform services are enabled by default.

To enable platform services:

1. Connect to the switch and log in as admin.
2. Issue the `msplmgmtactivate` command.

Example:

```
switch:admin> msplmgmtactivate
Request to activate MS Platform Service in progress.....
*Completed activating MS Platform Service in the fabric!
switch:admin>
```

To disable platform services:

1. Connect to the switch and log in as admin.
2. Issue the `msplmgmtdeactivate` command.
3. Enter `y` to confirm the deactivation.

Example:

```
switch:admin> msplmgmtdeactivate
MS Platform Service is currently enabled.
This will erase MS Platform Service configuration
information as well as database in the entire fabric.
Would you like to continue this operation? (yes, y, no, n): [no] y
Request to deactivate MS Platform Service in progress.....
*Completed deactivating MS Platform Service in the fabric!
switch:admin>
```

## Controlling access

You can use the `msconfigure` command to control access to the management server database.

An access control list (ACL) of WWN addresses determines which systems have access to the management server database. The ACL typically contains those WWNs of host systems that are running management applications.

If the list is empty (the default), the management server is accessible to all systems connected in-band to the fabric. For more access security, you can specify WWNs in the ACL so that access to the management server is restricted to only those WWNs listed.

The ACL is switch-based. Therefore, only hosts that are connected directly to the switch are affected by the ACL. A host that is somewhere else in the fabric and is connected to a switch with an empty ACL is allowed to access the management server.



**NOTE:** The `msconfigure` command is disabled if the switch is in secure mode. Refer to the *HP StorageWorks Secure Fabric OS user guide* for more information.

To display the management server ACL:

1. Connect to the switch and log in as admin.
2. Issue the `msconfigure` command.

The command becomes interactive.

3. At the select prompt, enter 1 to display the access list.

A list of WWNs that have access to the management server is displayed.

In this example, the list is empty:

```
switch:admin> msconfigure
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 1
MS Access list is empty.
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
done ...
switch:admin>
```

To add a member to the ACL:

1. Connect to the switch and log in as admin.
2. Issue the `msconfigure` command.  
The command becomes interactive.
3. At the select prompt, enter 2 to add a member based on its port/node WWN.
4. Enter the WWN of the host to be added to the ACL.
5. At the prompt, enter 1 to verify the WWN you entered was added to the ACL.
6. After verifying that the WWN was added correctly, enter 0 at the prompt to end the session.
7. At the `Update the FLASH?` prompt, enter `y`.
8. Press **Enter** to update the nonvolatile memory and end the session.

**Example:**

```
switch:admin> msconfigure
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 2
Port/Node WWN (in hex): [00:00:00:00:00:00:00:00] 20:00:00:20:37:65:ce:aa
*WWN is successfully added to the MS ACL.
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [2] 1
MS Access List consists of (14): {
20:00:00:20:37:65:ce:aa
20:00:00:20:37:65:ce:bb
20:00:00:20:37:65:ce:ff
20:00:00:20:37:65:ce:11
20:00:00:20:37:65:ce:22
20:00:00:20:37:65:ce:33
20:00:00:20:37:65:ce:44
10:00:00:60:69:04:11:24
10:00:00:60:69:04:11:23
21:00:00:e0:8b:04:70:3b
10:00:00:60:69:04:11:33
20:00:00:20:37:65:ce:55
20:00:00:20:37:65:ce:66
00:00:00:00:00:00:00:00
}
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 0
done ...
Update the FLASH? (yes, y, no, n): [yes] y
*Successfully saved the MS ACL to the flash.
switch:admin>
```

To delete a member from the ACL:

1. Connect to the switch and log in as admin.
2. Issue the **msconfigure** command.

The command becomes interactive.

3. At the select prompt, enter 3 to delete a member based on its port/node WWN.
4. At the prompt, enter the WWN of the member to be deleted from the ACL.
5. At the prompt, enter 1 to verify the WWN you entered was deleted from the ACL.
6. After verifying that the WWN was deleted correctly, enter 0 at the prompt to end the session.
7. At the Update the FLASH? prompt, enter **y**.
8. Press **Enter** to update the nonvolatile memory and end the session.

Example:

```
switch:admin> msconfigure
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 3
Port/Node WWN (in hex): [00:00:00:00:00:00:00] 20:00:00:20:37:65:ce:aa
*WWN is successfully deleted from the MS ACL.
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [2] 1
MS Access List consists of (13): {
20:00:00:20:37:65:ce:aa
20:00:00:20:37:65:ce:bb
20:00:00:20:37:65:ce:ff
20:00:00:20:37:65:ce:11
20:00:00:20:37:65:ce:22
20:00:00:20:37:65:ce:33
10:00:00:60:69:04:11:24
10:00:00:60:69:04:11:23
21:00:00:e0:8b:04:70:3b
10:00:00:60:69:04:11:33
20:00:00:20:37:65:ce:55
20:00:00:20:37:65:ce:66
}
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 0
done ...
Update the FLASH? (yes, y, no, n): [yes] y
*Successfully saved the MS ACL to the flash.
switch:admin>
```

---

## Configuring the server database

The management server database can be viewed or cleared.

To view the contents of the management server database:

1. Connect to the switch and log in as admin.
2. Issue the `msplatshow` command.

The contents of the management server platform database are displayed.

Example:

```
switch:admin> msplatshow
-----
Platform Name: [9] "first obj"
Platform Type: 5 : GATEWAY
Number of Associated M.A.: 1
[35] "http://java.sun.com/products/plugin"
Number of Associated Node Names: 1
Associated Node Names:
10:00:00:60:69:20:15:71
-----
Platform Name: [10] "second obj"
Platform Type: 7 : HOST_BUS_ADAPTER
Number of Associated M.A.: 1
Associated Management Addresses:
[30] "http://java.sun.com/products/1"
Number of Associated Node Names: 1
Associated Node Names:
10:00:00:60:69:20:15:75
```

To clear the management server database:

1. Connect to the switch and log in as admin.
2. Issue the `msplcleardb` command.
3. Enter `y` to confirm the deletion.

The management server platform database is cleared.

---

## Controlling topology discovery

The topology discovery feature can be displayed, enabled, and disabled; it is disabled by default.

To display topology discovery status:

1. Connect to the switch and log in as admin.
2. Issue the `mstdreadconfig` command.

Example:

```
switch:admin> mstdreadconfig
*MS Topology Discovery is Enabled.
switch:admin>
```

To enable topology discovery:

1. Connect to the switch and log in as admin.
2. Issue the `mstdenable` command to enable the discovery feature locally.
3. Issue the `mstdenable all` command to enable the discovery feature on the entire fabric.

Example:

```
switch:admin> mstdenable
Request to enable MS Topology Discovery Service in progress....
*MS Topology Discovery enabled locally.

switch:admin> mstdenable ALL
Request to enable MS Topology Discovery Service in progress....
*MS Topology Discovery enabled locally.
*MS Topology Discovery Enable Operation Complete!!
```

To disable topology discovery:

1. Connect to the switch and log in as admin.
2. Issue the `mstddisable` command to disable the discovery feature locally.  
A warning that all NID entries might be cleared is displayed.
3. Enter `y` to disable the discovery feature.
4. Issue the `mstddisable all` command to disable the discovery feature on the entire fabric.
5. Enter `y` to disable the discovery feature.



**NOTE:** Disabling management server topology discovery may erase all NID entries.

Example:

```
switch:admin> mstddisable
This may erase all NID entries. Are you sure? (yes, y, no, n): [no] y

Request to disable MS Topology Discovery Service in progress....
*MS Topology Discovery disabled locally.

switch:admin> mstddisable all
This may erase all NID entries. Are you sure? (yes, y, no, n): [no] y

Request to disable MS Topology Discovery Service in progress....
*MS Topology Discovery disabled locally.
*MS Topology Discovery Disable Operation Complete!!
```



# 12 Working with diagnostic features

This chapter contains the following sections:

- [Viewing power-on self test](#), page 165
- [Viewing switch status](#), page 166
- [Viewing port information](#), page 168
- [Viewing equipment status](#), page 171
- [Viewing the system message log](#), page 172
- [Viewing the port log](#), page 173
- [Configuring for syslogd](#), page 175
- [Viewing and saving diagnostic information](#), page 177
- [Setting up automatic trace dump transfers](#), page 178

This chapter provides information on diagnostics and how to display system, port, and specific hardware information. It also describes how to set up system logging mapping (`syslogd`) and how to set up the offloading of error messages (`supportsave`).

The purpose of the diagnostic subsystem is to evaluate the integrity of the system hardware.

Diagnostics are invoked two ways:

- Automatically during the power-on self test (POST)
- Manually using Fabric OS CLI commands

The error messages generated during these test activities are sent to the serial console and system message logs, whose output formats may differ slightly.

Use the `diaghelp` command to receive a list of all available diagnostic commands.

Refer to the *HP StorageWorks Fabric OS 4.x command reference guide* for a complete description of each command.

## Viewing power-on self test

By default, when you power on the system, the boot loader automatically performs power-on self tests and loads a Fabric OS kernel image.

The POST provides a quick indication of hardware readiness when hardware is powered up. These tests do not require user input to function. They typically operate within several minutes, and support minimal validation because of the restriction on test duration. Their purpose is to give a basic health check before a new switch joins a fabric.

These tests are divided into two groups: POST1 and POST2. POST1 validates the hardware interconnect of the device, and POST2 validates the ability of the device to pass data frames between the ports. The specific set of diagnostic and test commands run during POST depends on the switch model.

POST1 cannot be bypassed; it runs from the boot loader. The factory default configuration is also set to run POST2, but you can configure your switch to bypass POST2, which runs after the kernel image has started but before general system services, such as login, are enabled.

Although each test performed during POST2 is configurable, modify a POST2 test only if directed by your switch provider's customer service representative.

You can use the `diagdisablepost` command to disable POST2, and you can reenable it using the `diagenablepost` command. Refer to the *HP StorageWorks Fabric OS 4.x command reference guide* for additional information about these commands.

The following example shows a typical boot sequence, including POST messages:

```
*****
***  
* Copyright (c) 2004, *  
* Switchname *  
* Firmware Version 4.x.x.x.H build #XX: Sat Jan 31 12:40:41 PDT 2004  
* MAC Address: 00:05:1E:31:2F:60  
* BCSR FPGA Version 2007  
*  
*****  
***  
  
Power-on Self Test  
  
1024 MB SDRAM installed  
Executing code from SDRAM at 0203B5E0  
Power supply 1 status test PASSED  
Power supply 2 status test FAILED  
Fan tray status 1 test PASSED  
Fan tray status 2 test PASSED  
SEEPROM initialized. SEEPROM test is BYPASSED  
Dram test #2 is BYPASSED  
Compact Flash is installed  
  
Autoboot command: "memboot"  
Press <Enter> to execute or any other key to abort.
```

If you choose to bypass POST2, or after POST2 completes, various system services are started and the boot process displays additional console status and progress messages.

## Viewing switch status

Use the `switchstatusshow` command to display the overall status of the switch, including its power supplies, fans, and temperature. If the status of any one of these components is either marginal or down, the overall status of the switch is also displayed as marginal or down. If all components have a healthy status, the switch displays a healthy status.

The rules used to classify the health of each component are in the configuration files. To change the rules, you must use the `configupload` command to copy the configuration files, edit them manually, and use the `configdownload` command to load them back to the switch.

To view the overall status of the switch:

1. Connect to the switch and log in as admin.

**2.** Issue the `switchstatusshow` command:

```
switch:admin> switchstatusshow
[7505]: Read 1 license entries for generation 1.
[7505]: Read 1 license records.
Switch Health Report                               Report time: 05/21/2004
03:50:36 PM
Switch Name:      SW3900
IP address:      10.33.54.176
SwitchState:     MARGINAL
Duration:        863:23
Power supplies monitor  MARGINAL
Temperatures monitor   HEALTHY
Fans monitor          HEALTHY
Flash monitor         HEALTHY
Marginal ports monitor  HEALTHY
Faulty ports monitor   HEALTHY
Missing SFPs monitor    HEALTHY
All ports are healthy
switch:admin>
```

For more information on how the overall switch status is determined, refer to the `switchstatuspolicyset` command in the *HP StorageWorks Fabric OS 4.x command reference guide*.

To display switch information:

- 1.** Connect to the switch and log in as admin.
- 2.** Issue the `switchshow` command. This command displays the following information for a switch:
  - `switchname` displays the switch name.
  - `switchtype` displays the switch model and firmware version numbers.
  - `switchstate` displays the switch state: Online, Offline, Testing, or Faulty.
  - `switchrole` displays the switch role: Principal, Subordinate, or Disabled.
  - `switchdomain` displays the switch Domain ID.
  - `switchid` displays the embedded port D\_ID of the switch.
  - `switchwwn` displays the switch World Wide Name.
  - `switchbeacon` displays the switch beaconing state: either ON or OFF.

The `switchshow` command also displays the following information for ports on the specified switch:

- **Module type:** The SFP type if an SFP is present.
- **Port speed:** The speed of the Port (1G, 2G, 4G, N1, N2, N4, or AN). The speed can be fixed, negotiated, or auto negotiated.
- **Port state:** The port status.
- **Comment:** Displays information about the port. This section might be blank or display the WWN for the F\_Port or E\_Port, trunking state upstream or downstream status.

The details displayed for each switch differ on different switch models. For more information refer to the `switchshow` command in the *HP StorageWorks Fabric OS 4.x command reference guide*.

To display the uptime for a switch:

1. Connect to the switch and log in as admin.
2. Issue the `uptime` command:

```
switch:admin> uptime
4:43am  up 1 day, 12:32,  1 user,  load average: 1.29, 1.31, 1.27
switch:admin>
```

The `uptime` command displays the length of time the system has been in operation, the total cumulative amount of uptime since the system was first powered-on, the date and time of the last reboot, the reason for the last reboot, and the load average over the past one minute (1.29 in the preceding example), five minutes (1.31 in the example), and 15 minutes (1.27 in the example). The reason for the last switch reboot is also recorded in the system message log.

---

## Viewing port information

Use the commands that follow to view information about ports.

To view the status of a port:

1. Connect to the switch and log in as admin.
2. Issue the `portshow` command, specifying the number that corresponds to the port you are troubleshooting. In this example, the status of port 2 is shown:

```
switch:admin> portshow 2
      port 2 info
      Configuration   Current
Name :          port 2
State:          STARTED        UP
Type :          FC             FC
Link Status:    ENABLED        DOWN
Topology:       P-P            P-P
Speed:          AN             AN
LinkCost:       AUTO
WWN:            20:02:00:05:1e:12:f2:00

      Licensed     : YES
      Diag result  : PASSED

      inFrames:    0
      outFrames:   0
      inOctets:   0
      outOctets:  0
      discards:   0

switch:admin>
```

Refer to the *HP StorageWorks Fabric OS 4.x command reference guide* for additional `portshow` command information, such as the syntax for slot or port numbering.

To display the port statistics:

1. Connect to the switch and log in as admin.
2. Issue the `portstatsshow` command.

Port statistics include information such as number of frames received, number of frames sent, number of encoding errors received, and number of class 2 and class 3 frames received.

Refer to the *HP StorageWorks Fabric OS 4.x command reference guide* for additional `portstatsshow` command information, such as the syntax for slot or port numbering.

Example:

```
switch:admin> portstatsshow 3/7
stat_wtx          0        4-byte words transmitted
stat_wrx          0        4-byte words received
stat_ftx          0        Frames transmitted
stat_frx          0        Frames received
stat_c2_frx       0        Class 2 frames received
stat_c3_frx       0        Class 3 frames received
stat_lc_rx        0        Link control frames received
stat_mc_rx        0        Multicast frames received
stat_mc_to        0        Multicast timeouts
stat_mc_tx        0        Multicast frames transmitted
tim_rdy_pri       0        Time R_RDY high priority
tim_txcrd_z       0        Time BB_credit zero
er_enc_in         0        Encoding errors inside of frames
er_crc            0        Frames with CRC errors
er_trunc           0        Frames shorter than minimum
er_toolong          0        Frames longer than maximum
er_bad_eof          0        Frames with bad end-of-frame
er_enc_out          0        Encoding error outside of frames
er_disc_c3          0        Class 3 frames discarded
open               0        loop_open
transfer            0        loop_transfer
opened              0        FL_Port opened
starve_stop          0        tenancies stopped due to starvation
fl_tenancy           0        number of times FL has the tenancy
nl_tenancy           0        number of times NL has the tenancy
switch:admin>
```

To display a summary of port errors for a switch:

1. Connect to the switch and log in as admin.
2. Issue the `porterrshow` command. Refer to the *HP StorageWorks Fabric OS 4.x command reference guide* for additional `porterrshow` command information

Example:

```
switch:admin> porterrshow
          frames enc crc too bad enc disc link loss loss frjt fbsy
          tx   rx   in  err shrt long eof  out   c3 fail sync sig
sig=====
=
 0:  22  24    0    0    0    0    0 1.5m  0    7    3    0    0    0
 1:  22  24    0    0    0    0    0 1.2m  0    7    3    0    0    0
 2:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 3:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 4: 149m 99m    0    0    0    0    0    448   0    7    6    0    0    0
 5: 149m 99m    0    0    0    0    0    395   0    7    6    0    0    0
 6: 147m 99m    0    0    0    0    0    706   0    7    6    0    0    0
 7: 150m 99m    0    0    0    0    0    160   0    7    5    0    0    0
 8:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 9:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
10:   0    0    0    0    0    0    0    0    0    0    0    0    0    0
11:   0    0    0    0    0    0    0    0    0    0    0    0    2    0    0
12:   0    0    0    0    0    0    0    0    0    0    0    0    2    0    0
13:   0    0    0    0    0    0    0    0    0    0    0    0    2    0    0
14:   0    0    0    0    0    0    0    0    0    0    0    0    2    0    0
15:   0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
32:   0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
33:   0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
34:   0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
35:   0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
36:   0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
37:   0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
38:   0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
39:   0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
40: 99m 146m    0    0    0    0    0    666   0    6    796   7    0    0
41: 99m 149m    0    0    0    0    0    15k   0    2    303   4    0    0
42: 99m 152m    0    0    0    0    0    665   0    2    221   5    0    0
43: 99m 147m    0    0    0    0    0    16k   0    2    144   4    0    0
44:    0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
45:    0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
46:    0    0    0    0    0    0    0    0    0    0    0    0    2    0    0
47:    0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
```

The porterrshow command output provides one output line per port. See [Table 26](#) for a description of the error types.

**Table 26** Error summary description

Error type	Description
frames tx	Frames transmitted.
frames rx	Frames received.
enc in	Encoding errors inside frames.
crc err	Frames with CRC errors.
too shrt	Frames shorter than minimum.
too long	Frames longer than maximum.

**Table 26** Error summary description (continued)

Error type	Description
bad eof	Frames with bad end-of-frame delimiters.
enc out	Encoding error outside of frames.
disc c3	Class 3 frames discarded.
link fail	Link failures (LF1 or LF2 states).
loss sync	Loss of synchronization.
loss sig	Loss of signal.
frjt	Frames rejected with F_RJT.
fbsy	Frames busied with F_BSY.

## Viewing equipment status

You can display status for fans, power supply, and temperature.



**NOTE:** The number of fans, power supply units, and temperature sensors depends on the switch type. For detailed specifications on these components, refer to the switch hardware reference manual. The output from the status commands varies depending on the switch type

To display the status of the fans:

1. Connect to the switch and log in as admin.
2. Issue the `fanshow` command:

```
switch:admin> fanshow
Fan 1  Status: OK    Set_Speed: NORMAL      Actual_Speed: 7010 RPM
Fan 2  Status: OK    Set_Speed: NORMAL      Actual_Speed: 7180 RPM
Fan 3  Status: OK    Set_Speed: NORMAL      Actual_Speed: 7068 RPM
Fan 4  Status: OK    Set_Speed: NORMAL      Actual_Speed: 7116 RPM
Fan 5  Status: OK    Set_Speed: NORMAL      Actual_Speed: 7155 RPM
Fan 6  Status: OK    Set_Speed: NORMAL      Actual_Speed: 7001 RPM
switch:admin>
```

The possible status values are:

- OK Fan is functioning correctly.
- Absent Fan is not present.
- Below minimum Fan is present but rotating too slowly or stopped.
- Above minimum Fan is rotating too fast.
- Unknown Unknown fan unit installed.
- FAULTY Fan has exceeded hardware tolerance.

To display the status of a power supply:

1. Connect to the switch and log in as admin.
2. Issue the `psshow` command:

```
switch:admin> psshow
Power Supply #1 is OK
0335,FF2Z0007161,60-0000739-02, B,,DCJ3002-01P, B,FF2Z0007161
Power Supply #2 is faulty
0335,FF2Z0007176,60-0000739-02, B,,DCJ3002-01P, B,FF2Z0007176
switch:admin>
```

The possible status values are:

- OK Power supply functioning correctly.
- Absent Power supply not present.
- Unknown Unknown power supply unit installed.
- Predicting failure Power supply is present but predicting failure.
- Unknown Unknown fan unit installed.
- FAULTY Power supply present, but faulty (no power cable, power switch turned off, fuse blown, or other internal error).

To display temperature status:

1. Connect to the switch and log in as admin.
2. Issue the `tempshow` command:

```
switch:admin> tempshow

Index      Status      Centigrade      Fahrenheit
-----
1          OK           21              70
2          OK           22              72
3          OK           29              84
4          OK           24              75
5          OK           25              77
switch:admin>
```

Information displays for each temperature sensor in the switch.

The possible temperature status values are:

- OK Temperature is within acceptable range.
- FAIL Temperature is outside acceptable range. Damage may occur.

## Viewing the system message log

The system message log feature enables messages to be saved across power cycles and reboots.

The Core Switch 2/64 and the SAN Director 2/128 maintain an independent system message log for each of the two CP cards. For these models, you should configure `syslogd` to support chronological system message logs. For details, see “[Configuring for `syslogd`](#)” on page 175.

For details on error messages, refer to the *HP StorageWorks Fabric OS 4.x diagnostics and system error messages reference guide*.

To display the system message log, with no page breaks:

1. Connect to the switch and log in as admin.
2. Issue the `errdump` command.

To display the system message log, with page breaks:

1. Connect to the switch and log in as admin.
2. Issue the `errshow` command.

To clear the system message log:

1. Connect to the switch and log in as admin.
2. Issue the `errclear` command.

All switch and chassis events are removed.

---

## Viewing the port log

The Fabric OS maintains an internal log of all port activity. The port log stores entries for each port as a circular buffer. Each port has space to store 8000 log entries. When the log is full, the newest log entries overwrite the oldest log entries. Port logs are not persistent and are lost over power-cycles and reboots. If the port log is disabled, an error message displays.



**NOTE:** Port log functionality is completely separate from the system message log. Port logs are typically used to troubleshoot device connections.

---

To view the port log:

1. Connect to the switch and log in as admin.

**2. Issue the portlogshow command:**

```
switch:admin> portlogshow 8
Total records present      = 12
Number of records displayed = 12

Time        Module    Event     Port Len Log info
-----
18:36:52.036 fabctl   PrtSCN    08  0   st=1, Topo=2, Spd=0
18:36:52.361 WKA      Rx       08  140
22fffffe,00000000,01a6ffff,04000000
18:36:52.362 fabctl   PrtSCN    08  0   st=2, Topo=2, Spd=2
18:36:52.365 fabctl   Debug     08  0   Loading routes
18:36:52.379 fabctl   Tx       08  140
23640800,00fffffe,01a60001,02000000
18:36:52.379 WKA      Rx       08  140
22fffffc,00640800,02ceffff,03000000
18:36:52.382 nsd      Tx       08  140
23640800,00fffffc,02ceffff,02000000
18:36:52.382 WKA      Rx       08  32
22fffffd,00640800,02cdffff,62000000
18:36:52.383 fabctl   Tx       08  28
23640800,00fffffd,02cd0001,02000000
18:36:52.383 WKA      Ct_in    08  52
02fffffc,00640800,02d1ffff,01000000
18:36:52.384 nsd      Tx       08  40
03640800,00fffffc,02d1ffff,01000000
18:36:52.384 WKA      Ct_in    08  84
02fffffc,00640800,02d0ffff,01000000
```

Use the commands summarized in [Table 27](#) to view and manage port logs.

**Table 27** Commands for port log management

Command	Description
portlogclear	Clear port logs for all or particular ports.
portlogdisable	Disable port logs for all or particular ports.
portlogdump	Display port logs for all or particular ports, without page breaks.
portlogenable	Enable port logs for all or particular ports.
portlogshow	Display port logs for all or particular ports, with page breaks.

NOTE: Refer to the *HP StorageWorks Fabric OS 4.x command reference guide* for additional information about these commands.

The portlogdump command output (trace) is a powerful tool that is used to troubleshoot fabric issues. The portlogdump output provides detailed information about the actions and communications within a fabric. By understanding the processes that are taking place in the fabric, issues can be identified and located.

The portlogdump command displays the port log, showing a portion of the Fibre Channel payload and header (FC-PH). The header contains control and addressing information associated with the frame. The payload contains the information being transported by the frame and is determined by the higher-level service or FC\_4 upper level protocol. There are many different payload formats based on the protocol.

Because a portlogdump output is long, a truncated example is presented:

```
switch:admin> portlogdump
task event port cmd args
-----
16:30:41.780 PORT Rx 9 40 02fffffd,00fffffd,0061ffff,14000000
16:30:41.780 PORT Tx 9 0 c0fffffd,00fffffd,0061030f
16:30:42.503 PORT Tx 9 40 02fffffd,00fffffd,0310ffff,14000000
16:30:42.505 PORT Rx 9 0 c0fffffd,00fffffd,03100062
16:31:00.464 PORT Rx 9 20 02fffc01,00fffc01,0063ffff,01000000
16:31:00.464 PORT Tx 9 0 c0fffc01,00fffc01,00630311
16:31:00.465 nsd ctn 9 fc 000104a0,0000007f
16:31:00.465 nsd cout 9 fc 00038002,00000003,01fffc01
16:31:00.466 PORT Tx 9 356 03fffc01,00fffc01,00630311,01000000
16:31:00.474 PORT Rx 9 0 c0fffc01,00fffc01,00630311
16:31:01.844 PORT Tx 9 40 02fffffd,00fffffd,0312ffff,14000000
16:31:01.854 PORT Rx 9 0 c0fffffd,00fffffd,03120064
16:31:01.963 PORT Rx 9 40 02fffffd,00fffffd,0065ffff,14000000
16:31:01.963 PORT Tx 9 0 c0fffffd,00fffffd,00650313
16:31:14.726 INTR pstate 0 LF2
16:31:14.729 PORT scn 0 137 00000000,00000000,00000008
16:31:14.729 PORT scn 0 129 00000000,00000000,00000400
16:31:14.729 PORT scn 0 2 00010004,00000000,00000002
16:31:14.730 SPEE sn 0 ws 00000002,00000000,00000000
<output truncated>
```

Refer to the *HP StorageWorks Fabric OS 4.x command reference guide* for information on the portlogdump command.

## Configuring for syslogd

The system logging daemon (syslogd) is a process on UNIX, Linux, and some Windows systems that reads and logs messages as specified by the system administrator.

Fabric OS can be configured to use a UNIX-style syslogd process to forward system events and error messages to log files on a remote host system.

The host system can be running UNIX, Linux, or any other operating system that supports the standard syslog functionality.



**NOTE:** Fabric OS releases earlier than v4.4.0 did not support UNIX local7 facilities; they supported kern facilities. Starting with Fabric OS v4.4.0, kern facilities are no longer supported; UNIX local7 facilities are supported (the default facility level is 7).

Configuring for syslogd involves configuring the host, enabling syslogd on the HP StorageWorks switch, and, optionally, setting the facility level.

# Configuring the Host

Fabric OS supports a subset of UNIX-style message severities that default to the UNIX local7 facility. To configure the host, edit the `/etc/syslog.conf` file to map Fabric OS message severities to UNIX severities, as shown in [Table 28](#).

**Table 28** Fabric OS to UNIX message severities

Fabric OS message severity	UNIX message severity
Critical (1)	Emergency (0)
Error (2)	Error (3)
Warning (3)	Warning (4)
Info (4)	Info (6)

In the following example, Fabric OS messages map to local7 facility level 7 in the `/etc/syslog.conf` file:

```
local7.emerg      /var/adm/swcritical
local7.alert      /var/adm/alert7
local7.crit       /var/adm/crit7
local7.err        /var/adm/swerror
local7.warning    /var/adm/swwarning
local7.notice     /var/adm/notice7
local7.info       /var/adm/swinfo
local7.debug      /var/adm/debug7
```

If you prefer to map Fabric OS severities to a different UNIX local7 facility level, see "[To set the facility level](#):" on page 177.

# Configuring the switch

Configuring the switch involves specifying syslogd hosts and, optionally, setting the facility level. You can also remove a host from the list of syslogd hosts.

To specify syslogd hosts:

1. Connect to the switch and log in as admin.
2. Issue the `syslogdipadd` command and specify an IP address.
3. Verify the IP address was entered correctly using the `syslogdipshow` command.

You can specify up to six host IP addresses for storing syslog messages, as shown in this example:

```
switch:admin> syslogipadd 10.1.2.1
switch:admin> syslogipadd 10.1.2.2
switch:admin> syslogipadd 10.1.2.3
switch:admin> syslogipadd 10.1.2.4
switch:admin> syslogipadd 10.1.2.5
switch:admin> syslogipadd 10.1.2.6
switch:admin> syslogipshow
syslog.IP.address.1 10.1.2.1
syslog.IP.address.2 10.1.2.2
syslog.IP.address.3 10.1.2.3
syslog.IP.address.4 10.1.2.4
syslog.IP.address.5 10.1.2.5
syslog.IP.address.6 10.1.2.6
```

To set the facility level:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
syslogdfacility -1 n
```

where *n* is a number from 0 through 7, indicating a UNIX local7 facility. The default is 7.

It is necessary to set the facility level only if you specified a facility other than local7 in the host /etc/syslog.conf file.

To remove a syslogd host from the list:

1. Connect to the switch and log in as admin.
2. Issue the **syslogipremove** command:
3. Verify the IP address was deleted using the **syslogipshow** command.

```
switch:admin> syslogipremove 10.1.2.1
```

---

## Viewing and saving diagnostic information

Issue the **supportshow** command to dump important diagnostic and status information to the session screen, where you can review it or capture its data.

To save a set of files that customer support technicians can use to further diagnose the switch condition, issue the **supportsave** command. The command prompts for an FTP server, packages the following files, and sends them to the specified server:

- The output of the **supportshow** command
- The contents of any trace dump files on the switch
- System message logs (for the Core Switch 2/64 and the SAN Director 2/128, **supportsave** saves the system message logs from both of the CP cards)

# Setting up automatic trace dump transfers

You can set up a switch so that diagnostic information is transferred automatically to a remote server. Then, if a problem occurs you can provide your customer support representative with the most detailed information possible. To ensure the best service, you should set up for automatic transfer as part of standard switch configuration, before a problem occurs.

Setting up for automatic transfer of diagnostic files involves the following tasks:

- Specify a remote server to store the files.
- Enable the automatic transfer of trace dumps to the server. (Trace dumps overwrite each other by default; sending them to a server preserves information that would otherwise be lost.)
- You should also set up a periodic checking of the remote server so that you are alerted if the server becomes unavailable and you can correct the problem.

Once the setup is complete, you can run the `supportsave -c` command to save diagnostic information to the server without the need to specify server details.

The following procedures describe in detail the tasks for setting up automatic transfer. For details on the commands, refer to the *HP StorageWorks Fabric OS 4.x command reference guide*.

To specify a remote server:

1. Verify that the FTP service is running on the remote server.
2. Connect to the switch and log in as admin.
3. Issue the following command:

```
supportftp -s
```

The command becomes interactive and you are prompted for the required information.

4. Respond to the prompts as follows:

• Host Name	Enter the name or IP address of the server where the file is to be stored; for example, 192.1.2.3.
• User name	Enter the user name of your account on the server; for example, JohnDoe.
• Password	Enter your account password for the server.
• Remote directory	Specify a path name for the remote directory. Absolute path names can be specified using forward slash (/). Relative path names create the directory in the user's home directory on UNIX servers, and in the directory where the FTP server is running on Windows servers.

To enable the automatic transfer of trace dumps:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
traceftp -e
```

To set up periodic checking of the remote server:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
supportftp -t interval
```

The *interval* is in hours. The minimum interval is 1 hour. Specify 0 hours to disable the checking feature.

To save a comprehensive set of diagnostic files to the server:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
supportsave -c
```



# 13 Troubleshooting

This chapter contains the following sections:

- [Most common problem areas](#), page 182
- [Gathering information for technical support](#), page 182
- [Analyzing connection problems](#), page 184
- [Restoring a segmented fabric](#), page 186
- [Correcting zoning setup issues](#), page 187
- [Recognizing MQ-WRITE errors](#), page 189
- [Correcting I2C bus errors](#), page 190
- [Correcting device login issues](#), page 191
- [Identifying media-related issues](#), page 194
- [Correcting link failures](#), page 196
- [Correcting marginal links](#), page 198
- [Inaccurate information in the system message log](#), page 200
- [Recognizing the port initialization and FCP auto discovery process](#), page 200

Troubleshooting should begin at the center of the SAN — the fabric. Because switches are located between the hosts and storage devices, and have visibility into both sides of the storage network, starting with them can help narrow the search path. After eliminating the possibility of a fault within the fabric, see if the problem is on the storage side or the host side, and continue a more detailed diagnosis from there. Using this approach can quickly pinpoint and isolate problems.

For example, if a host cannot see a storage device, run a switch command to see if the storage device is logically connected to the switch. If not, focus first on the storage side. Use storage diagnostic tools to better understand why it is not visible to the switch. When the storage can be seen from the switch, if the host still cannot see the storage device, then there is still a problem between the host and switch.

This chapter provides information on troubleshooting and the most common procedures used to diagnose and repair issues. It also includes specific troubleshooting scenarios as examples.

# Most common problem areas

See [Table 29](#) for a list of the most common problem areas that arise within SANs and a list of tools that can be used to resolve the problems.

**Table 29** Common troubleshooting problems and tools

Problem area	Investigate	Tools
Fabric	<ul style="list-style-type: none"><li>Missing devices</li><li>Marginal links (unstable connections)</li><li>Incorrect zoning configurations</li><li>Incorrect switch configurations</li></ul>	<ul style="list-style-type: none"><li>Switch LEDs</li><li>Switch commands for diagnostics (command line)</li><li>Web or GUI-based monitoring and management software tools</li><li>Real-time distributed fabric operating system with advanced diagnostics</li></ul>
Storage devices	<ul style="list-style-type: none"><li>Physical issues between switch and devices</li><li>Incorrect storage software configurations</li></ul>	<ul style="list-style-type: none"><li>Device LEDs</li><li>Storage diagnostic tools</li></ul>
Hosts	<ul style="list-style-type: none"><li>Incorrect host bus adapter installation</li><li>Incorrect device driver installation</li><li>Incorrect device driver configuration</li></ul>	<ul style="list-style-type: none"><li>Host adaptor LEDs</li><li>Host operating system diagnostic tools</li><li>Device driver diagnostic tools</li></ul>
Storage management applications	Incorrect installation and configuration of the storage devices that the software references. For example, if using a volume-management application, check for: <ul style="list-style-type: none"><li>Incorrect volume installation</li><li>Incorrect volume configuration</li></ul>	Application-specific tools and resources

# Gathering information for technical support

Common steps and questions to ask yourself when troubleshooting a system problem are as follows:

1. What is current Fabric OS level?
2. What is switch hardware version?
3. Is the switch operational?

- 4.** Impact assessment and urgency:
  - Is the switch down?
  - Is it a standalone switch?
  - How large is the fabric?
  - Is the fabric redundant?
- 5.** Issue the `supportSave` command. (See “[Viewing and saving diagnostic information](#)” on page 177 and “[Setting up automatic trace dump transfers](#)” on page 178).
- 6.** Document the sequence of events by answering the following questions:
  - What happened just prior to the problem?
  - Is the problem reproducible?
  - If so, what are the steps to produce the problem?
  - What configuration was in place when the problem occurred?
- 7.** Did a failover occur?
- 8.** Was security enabled?
- 9.** Was POST enabled?
- 10.** Are serial port (console) logs available?
- 11.** Which CP card was master? (applicable only to the Core Switch 2/64 and the SAN Director 2/128)
- 12.** What and when were the last actions or changes made to the system?

If the problem still cannot be resolved, retrieve as much of the following informational items as possible prior to contacting the SAN technical support vendor.

- 1.** Switch information:
  - Serial number (located on the chassis)
  - Worldwide name (obtain using `licenseIdShow` or `wwn` commands)
  - Fabric OS version (obtain using `version` command)
  - Switch Configuration settings
  - SupportSave output
- 2.** Host information:
  - OS version and patch level
  - HBA type
  - HBA firmware version
  - HBA driver version
  - Configuration settings

**3. Storage information:**

- Disk/tape type
- Disk/tape firmware level
- Controller type
- Controller firmware level
- Configuration settings
- Storage software (such as EMC Control Center, Veritas SPC, etc.)

---

## Analyzing connection problems

If a host is unable to detect its target (for example, a storage or tape device), you should begin troubleshooting the problem in the middle of the data path. Determine if the problem is *above* or *below* the starting point, then continue to divide the suspected problem path in half until you can pinpoint the problem.

Use the following procedures to analyze the problem:

To check the logical connection:

1. Issue the `switchShow` command.
2. Review the output and determine if the device is logically connected to the switch:
  - A device that *is* logically connected to the switch is registered as an `Nx_Port`.
  - A device that is *not* logically connected to the switch is registered as something other than an `Nx_Port`.
3. If the missing device *is* logically connected, proceed to the next troubleshooting procedure ("[To check the Simple Name Server \(SNS\):](#)" on page 185).
4. If the missing device is *not* logically connected, eliminate the host and everything on that side of the data path from the suspect list.

This includes all aspects of the host OS, the HBA driver settings and binaries, the HBA Basic Input Output System (BIOS) settings, the HBA SFP, the cable going from the switch to the host, the SFP on the switch side of that cable, and all switch settings related to the host. See "[To check for a link initialization failure \(loop\):](#)" on page 197 as the next potential trouble spot.

To check the Simple Name Server (SNS):

- Issue the `nsshow` command on the switch to which the device is attached.

```
The Local Name Server has 9 entries {

  Type   Pid    COS      PortName           NodeName          TTL(sec)

*N    021a00;    2,3;20:00:00:e0:69:f0:07:c6;10:00:00:e0:69:f0:07:c6; 895
      Fabric Port Name: 20:0a:00:60:69:10:8d:fd
NL   051edc;    3;21:00:00:20:37:d9:77:96;20:00:00:20:37:d9:77:96; na
      FC4s: FCP [SEAGATE ST318304FC      0005]

      Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL   051ee0;    3;21:00:00:20:37:d9:73:0f;20:00:00:20:37:d9:73:0f; na
      FC4s: FCP [SEAGATE ST318304FC      0005]

      Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL   051ee1;    3;21:00:00:20:37:d9:76:b3;20:00:00:20:37:d9:76:b3; na
      FC4s: FCP [SEAGATE ST318304FC      0005]

      Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL   051ee2;    3;21:00:00:20:37:d9:77:5a;20:00:00:20:37:d9:77:5a; na
      FC4s: FCP [SEAGATE ST318304FC      0005]

      Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL   051ee4;    3;21:00:00:20:37:d9:74:d7;20:00:00:20:37:d9:74:d7; na
      FC4s: FCP [SEAGATE ST318304FC      0005]

      Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL   051ee8;    3;21:00:00:20:37:d9:6f:eb;20:00:00:20:37:d9:6f:eb; na
      FC4s: FCP [SEAGATE ST318304FC      0005]

      Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL   051eef;    3;21:00:00:20:37:d9:77:45;20:00:00:20:37:d9:77:45; na
      FC4s: FCP [SEAGATE ST318304FC      0005]

      Fabric Port Name: 20:0e:00:60:69:10:9b:5b
N    051f00;    2,3;50:06:04:82:bc:01:9a:0c;50:06:04:82:bc:01:9a:0c; na
      FC4s: FCP [EMC      SYMMETRIX      5267]

      Fabric Port Name: 20:0f:00:60:69:10:9b:5b
```

- Look for the device in the SNS list, which lists the nodes connected to that switch, allowing you to determine if a particular node is accessible on the network.

- If the device is *not* present in the SNS, the problem is between the storage device and the switch. There might be a time-out or communication problem between edge devices and the Name Server. Check the edge device documentation to determine if there is a time-out setting or parameter that can be reconfigured. If this does not solve the communication problem, contact the support organization for the product that appears to be timing out.
- If the device *is* listed in the SNS, the problem is between the storage device and the host. There might be a zoning mismatch or a host/storage issue. Proceed to “[To check for zoning problems:](#)” on page 185.

To check for zoning problems:

- Issue the `cfgactvshow` command to determine if zoning is enabled.

If zoning is enabled, it is possible that the problem is being caused by zoning enforcement (for example, two devices in different zones cannot see each other).

- Confirm that the specific edge devices that need to communicate with each other are in the same zone.
  - If they are in the same zone, zoning is not causing the communication problem.
  - If they are not in the same zone and zoning is enabled, proceed to [step 3](#).

3. Resolve zoning conflicts by putting the devices into the same zoning configuration.

See “[Correcting zoning setup issues](#)” on page 187 for additional information.

---

## Restoring a segmented fabric

Fabric segmentation is generally caused by:

- Incompatible fabric parameters (see “[To reconcile fabric parameters individually:](#)” on page 186).
- Incorrect PID setting (see “[Configuring the PID format](#)” on page 203).
- Incompatible zoning configuration (see “[To check for zoning problems:](#)” on page 185).
- Domain ID conflict (see “[To reconcile a domain ID conflict:](#)” on page 187).
- A switch in a secure fabric is not running Secure Fabric OS.

Refer to the *HP StorageWorks Secure Fabric OS user guide* for additional information.

There are a number of settings that control the overall behavior and operation of the fabric. Some of these values, such as the domain ID, are assigned automatically by the fabric and can differ from one switch to another in the fabric. Other parameters, such as the BB credit, can be changed for specific applications or operating environments, but must be the same among all switches to allow the formation of a fabric.

The following fabric parameters must be identical for a fabric to merge:

- R\_A\_TOV
- E\_D\_TOV
- Data field size
- Sequence level switching
- Disable device probing
- Suppress class F traffic
- VC encoded address mode
- Perframe route priority
- long-distance fabric (not necessary on Bloom-based fabrics)
- BB credit
- Core PID

To reconcile fabric parameters individually:

1. Log in to one of the segmented switches as admin.
2. Issue the `configshow` command.
3. Open another telnet session and log in to another switch in the same fabric as admin.
4. Issue the `configshow` command.
5. Compare the two switch configurations line by line and look for differences. Do this by comparing the two telnet windows, or by printing the `configshow` output.
6. Connect to the segmented switch after the discrepancy is identified.
7. Disable the switch by entering the `switchdisable` command.

- Issue the `configure` command to edit the fabric parameters for the segmented switch.

Refer to the *HP StorageWorks Fabric OS 4.x command reference guide* for more detailed information.

- Enable the switch by entering the `switchenable` command.

To download a correct configuration:

You can restore a segmented fabric by downloading a previously saved correct backup configuration to the switch. Downloading in this manner reconciles any discrepancy in the fabric parameters and allows the segmented switch to rejoin the main fabric. For details on uploading and downloading configurations, see "[Maintaining configurations](#)" on page 75.

To reconcile a domain ID conflict:

If a domain ID conflict appears, the conflict is only reported at the point where the two fabrics are physically connected. However, there might be several conflicting domain IDs, which appear as soon as the initial conflict is resolved.

Repeat this procedures until all domain ID conflicts are resolved:

- Issue the `switchshow` command on a switch from one of the fabrics.
- Open a separate telnet window.
- Issue the `switchshow` command on a switch from the second fabric.
- Compare the `switchshow` output from the two fabrics. Note the number of domain ID conflicts; there might be several duplicate domain IDs that need to be changed.
- Choose the fabric on which to change the duplicate domain ID; connect to the conflicting switch in that fabric.
- Issue the `switchdisable` command.
- Issue the `switchenable` command.

This enables the joining switch to obtain a new domain ID as part of the process of coming online. The fabric principal switch allocates the next available domain ID to the new switch during this process.

- Repeat step 5 through step 7 if additional switches have conflicting domain IDs.

---

## Correcting zoning setup issues

The types of zone configuration discrepancies that can cause segmentation are listed in [Table 30](#).

**Table 30** Types of zone discrepancies

Conflict cause	Description
Configuration mismatch	Occurs when zoning is enabled in both fabrics and the zone configurations that are enabled are different in each fabric.
Type mismatch	Occurs when the name of a zone object in one fabric is also used for a different type of zone object in the other fabric. A zone object is any device in a zone.
Content mismatch	Occurs when the definition of a in one fabric is different from the definition of a zone object with the same name in the other fabric.

**Table 31** summarizes commands that are useful for debugging zoning issues.

**Table 31** Commands for debugging zoning

Command	Function
alicreate	Use to create a zone alias.
alidelete	Use to delete a zone alias.
cfgcreate	Use to create a zone configuration.
cfgshow	Displays zoning configuration.
licenseshow	Displays current license keys and associated (licensed) products.
switchshow	Displays currently enabled configuration and any E_Port segmentations due to zone conflicts.
zoneadd	Use to add a member to an existing zone.
zonecreate	Use to create a zone. Before a zone becomes active, the cfgSave and cfgEnable commands must be used.
zonehelp	Displays help information for zone commands.
zoneshow	Displays zone information.

See “[Administering advanced zoning](#)” on page 123 and the following documents for additional information about setting up zoning properly:

- The Advanced Zoning chapter in the *HP StorageWorks Fabric OS 4.x features overview guide*.
- The *Advanced Zoning Commands* section of the *HP StorageWorks Fabric OS 4.x command reference guide*.

To correct a fabric merge problem quickly:

You can correct zone conflicts by using the `cfgclear` command to clear the zoning database.



**CAUTION:** This is a disruptive procedure.

---

1. Determine which switches have the incorrect zoning configuration; then log in to the switches as admin.
2. Issue the `switchDisable` command.
3. Issue the `cfgClear` command.



**CAUTION:** This command clears the zoning database on the affected switches.

---

4. Issue the `cfgDisable` command.
5. Issue the `switchEnable` command. This forces a zone merge and populates the switches with the desired zoning database. The two fabrics are merged together again.

To correct a merge conflict without disrupting the fabric, first verify fabric merge problem, then edit zone configuration members, and then reorder the zone member list.

To verify a fabric merge problem:

1. Issue the `switchshow` command to validate that the segmentation is due to a zone issue.
2. See [Table 30](#) on page 187 to view the different types of zone discrepancies.

To edit zone configuration members:

1. Log in to one of the switches in a segmented fabric as admin.

2. Issue the `cfgshow` command.

Typing the asterisk \* after the command displays list of all configuration names.

3. Print the output from the `cfgShow` command.

4. Start another telnet session and connect to the next fabric as an administrator.

5. Issue the `cfgShow` command.

6. Print the output from the `cfgShow` command.

7. Compare the two fabric zone configurations line by line and look for an incompatible configuration.

8. Connect to one of the fabrics.

9. Issue zone configure edit commands to edit the fabric zone configuration for the segmented switch (see [Table 31](#) on page 188 for specific commands).

If the zoneset members between two switches are not listed in the same order in both configurations, the configurations are considered a mismatch; this results in the switches being segmented in the fabric.

For example:

`[cfg1 = z1; z2]` is different from `[cfg1 = z2; z1]`, even though the members of the configuration are the same.

One simple approach to making sure that the zoneset members are in the same order is to keep the members in alphabetical order.

To reorder the zone member list:

1. Use the output from the `cfgshow` for both switches.
2. Compare the order that the zone members are listed. Members must be listed in the same order.
3. Rearrange zone members so that the configuration for both switches is the same. Arrange zone members in alphabetical order, if possible.

---

## Recognizing MQ-WRITE errors

An MQ error is a message queue error. Identify an MQ error message by looking for the two letters M and Q in the error message.

Example:

```
2004/08/24-10:04:42, [MQ-1004], 218,, ERROR, ras007, mqRead, queue =
raslog-test-string0123456-raslog, queue I
D = 1, type = 2
```

MQ errors can result in devices dropping from the SNS or can prevent a switch from joining the fabric. MQ errors are rare and difficult to troubleshoot, and it is suggested that they be resolved by working with the switch supplier. When MQ errors are encountered, execute the `supportsave` command to capture debug information about the switch; then forward the `supportsave` data to the switch supplier for further investigation.

# Correcting I<sup>2</sup>C bus errors

I<sup>2</sup>C bus errors indicate defective hardware, and the specific item is listed in the error message. Refer to the *HP StorageWorks Fabric OS 4.x diagnostics and system error messages reference guide* for information specific to the error that was received. Some CPT and Environmental Monitor (EM) messages contain I<sup>2</sup>C-related information.

If the I<sup>2</sup>C message does not indicate the specific hardware that might be failing, begin debugging the hardware, as this is the most likely cause. The next sections provide procedures for debugging the hardware.

To check fan components:

1. Log in to the switch as user.
2. Issue the `fanshow` command.
3. Check the fan status and speed output.

If any of the fan speeds display abnormal RPMs, replace the fan FRU.

To check the switch temperature:

1. Log in to the switch as user.
2. Issue the `tempshow` command.
3. Check the temperature output.

Look for indications of high or low temperatures.

To check the power supply:

1. Log in to the switch as user.
2. Issue the `psshow` command.
3. Check the power supply status. Refer to the appropriate hardware reference manual for details regarding the power supply status.

If any of the power supplies show a status other than OK, consider replacing the power supply as soon as possible.

To check the temperature, fan, and power supply:

1. Log in to the switch as user.
2. Issue the `sensorshow` command. Refer to the *HP StorageWorks Fabric OS 4.x command reference guide* for details regarding the sensor numbers.
3. Check the temperature output.

Look for indications of high or low temperatures.

4. Check the fan speed output.

If any of the fan speeds display abnormal RPMs, replace the fan FRU.

5. Check the power supply status.

If any of the power supplies show a status other than OK, consider replacing the power supply as soon as possible.

# Correcting device login issues

To try to pinpoint problems with device logins, use this procedure:

1. Log in to the switch as root.
2. Issue the `switchshow` command; then check for correct logins. For example:

```
switch:admin> switchshow
switchName:      sw094135
switchType:      26.1
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:   126
switchId:        fffc7e
switchWwn:       10:00:00:05:1e:34:00:69
zoning:          ON (cfg_em)
switchBeacon:    OFF
Port   Media Speed State
=====
 0     id   N1   Online   E-Port  10:00:00:60:69:11:f9:fc
"2800_116"
 1     id   1G   Online   E-Port  10:00:00:60:69:11:f9:fc
"2800_116"
 2     id   N2   No_Light
 3     id   2G   No_Light
 4     id   N2   Online   E-Port  (Trunk port, master is Port  5)
 5     id   N2   Online   E-Port  10:00:00:05:1e:34:00:8b
"Dazz125" (downstream) (Trunk master)
 6     id   N2   No_Light
 7     id   N2   No_Light
 8     id   N1   Online   L-Port  4 public, 1 private, 1 phantom
 9     id   N2   No_Light
10    id   N2   Online   G-Port
11    id   N2   Online   F-Port  10:00:00:01:c9:28:c7:01
12    id   N1   Online   L-Port  4 public, 1 private, 1 phantom
13    --   N2   No_Module
14    id   N2   Online   E-Port  (Trunk port, master is Port  15)
15    id   N2   Online   E-Port  10:00:00:60:69:90:03:17
"TERM_113" (downstream) (Trunk master)
```

3. Issue the portconfigshow command to see how the port is configured. For example:

```
sw094135:root> portcfgshow
Ports of Slot 0    0  1  2  3    4  5  6  7    8  9 10 11   12 13 14 15
-----+-----+-----+-----+-----+-----+-----+-----+
-+-
Speed          AN 1G AN 2G    AN AN AN AN    AN AN AN AN    AN AN AN AN
Trunk Port     ON ON .. ON    ON ON ON ON    ON ON ON ON    ON ON ON
ON
Long Distance  ... ... ... ...
..
VC Link Init   ... ... ... ...
..
Locked L_Port   ... ... ... ...
..
Locked G_Port   ... ... ... ...
..
Disabled E_Port ... ... ... ...
..
ISL R_RDY Mode ... ... ... ...
..
Persistent Disable... ... ...
..
Locked Loop HD ... ... ...
..
p
where AN:AutoNegotiate, ..:OFF, ?:INVALID.
LM:L0.5
```

4. Issue the porterrshow command; then check for errors that can cause login problems.

```
sw094135:root> porterrshow
      frames enc  crc  too  too  bad  enc disc link loss loss frjt fbsy
      tx    rx  in   err shrt long eof  out   c3 fail sync sig
=====
0:  38   75   0   0   0   0   0   0   0   9   11   0   0   0
1: 110   73   0   0   0   0   0   0   0   9   11   0   0   0
2:   0   0   0   0   0   0   0   38   0   4   0   2   0   0
3:   0   0   0   0   0   0   0   0   0   4   1   2   0   0
4: 59m 102   0   0   0   0   0   0   0   4   0   0   0   0   0
5: 59m 103   0   0   0   0   0   0   0   3   0   0   0   0   0
6:   0   0   0   0   0   0   0   21   0   3   0   0   0   0
7:   0   0   0   0   0   0   0   58   0   3   0   0   0   0
8:  81   19k   0   0   0   0   0   0   3.0m  0   5   43   0   0
9:   0   0   0   0   0   0   0   29   0   3   0   0   0   0
10: 12m 68m   0   0   0   0   0   0   13   43m  8   1   1   0   0
11: 30m 33m   0   0   0   0   0   0   0   0   8   1   1   0   0
12:  89   25k   0   0   0   0   0   0   2.9m  0   7   43   0   0
13:   0   0   0   0   0   0   0   0   0   3   0   0   0   0
14: 29m 82m   0   0   0   0   0   0   0   1.2m  4   1   1   0   0
15: 29m 81m   0   0   0   0   0   0   0   1.1m  4   1   1   0   0
```

- A high number of errors relative to the frames transmitted and frames received can indicate a marginal link see “[Correcting marginal links](#)” on page 198 for additional information).
- A steadily increasing number of errors can indicate a problem. Track errors by sampling the port errors every five or ten seconds.

5. Issue the **portflagsshow** command; then check to see how a port has logged in and where a login failed (if a failure occurred).

```
sw094135:root> portflagsshow
Port SNMP      Physical   Flags
-----
  0: Online    In_Sync    PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
LOGIN LED
  1: Online    In_Sync    PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
LOGIN LED
  2: Offline   No_Light   PRESENT U_PORT LED
  3: Offline   No_Light   PRESENT U_PORT LED
  4: Online    In_Sync    PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
LOGIN LED
  5: Online    In_Sync    PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
LOGIN LED
  6: Offline   No_Light   PRESENT U_PORT LED
  7: Offline   No_Light   PRESENT U_PORT LED
  8: Online    In_Sync    PRESENT ACTIVE F_PORT L_PORT U_PORT LOGICAL_ONLINE
LOGIN NOELP LED ACCEPT
  9: Offline   No_Light   PRESENT U_PORT LED
10: Online    In_Sync    PRESENT ACTIVE G_PORT U_PORT LOGIN LED
11: Online    In_Sync    PRESENT ACTIVE F_PORT G_PORT U_PORT LOGICAL_ONLINE
LOGIN NOELP LED ACCEPT
12: Online    In_Sync    PRESENT ACTIVE F_PORT L_PORT U_PORT LOGICAL_ONLINE
LOGIN NOELP LED ACCEPT
13: Offline   No_Module  PRESENT U_PORT LED
14: Online    In_Sync    PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
LOGIN LED
15: Online    In_Sync    PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
LOGIN LED
```

6. Issue the **portlogdumpport portid** command; then view the device to switch communication.

```
sw094135:root> portlogdumpport 10
time          task       event     port cmd args
-----
12:38:21.590  SPEE      sn        10   WS  00000000,00000000,00000000
12:38:21.591  SPEE      sn        10   WS  000000ee,00000000,00000000
12:38:21.611  SPEE      sn        10   WS  00000001,00000000,00000000
12:38:21.871  SPEE      sn        10   NC  00000002,00000000,00000001
12:38:21.872  LOOP      loopscn  10   LIP  8002
12:38:22.171  LOOP      loopscn  10   TMO  2
12:38:22.171  INTR      pstate   10   LF2
12:38:22.172  INTR      pstate   10   OL2
12:38:22.172  INTR      pstate   10   LR3
12:38:22.172  INTR      pstate   10   AC
12:38:22.172  PORT      scn      10   11  00000000,00000000,00000002
12:38:22.311  PORT      scn      10   1   00000000,00000000,00000001
12:38:22.311  PORT      debug    10   00000001,00654320,00000001,00000000
12:38:22.311  PORT      debug    10   00000001,00654320,00000002,00000000
12:38:22.311  PORT      debug    10   00000001,00654320,00000003,00000000
12:38:22.313  PORT      Tx       10   164 02fffffd,00fffffd,025effff,10000000
12:38:22.314  PORT      debug    10   00000001,00654320,00000003,00000000 *
7
12:38:28.312  PORT      Tx       10   164 02fffffd,00fffffd,028fffff,10000000
12:38:34.312  PORT      Tx       10   164 02fffffd,00fffffd,0293ffff,10000000
12:38:40.312  PORT      Tx       10   164 02fffffd,00fffffd,0299ffff,10000000
12:38:46.312  PORT      Tx       10   164 02fffffd,00fffffd,029bffff,10000000
12:38:52.312  PORT      Tx       10   164 02fffffd,00fffffd,029dfffd,10000000
12:38:58.312  PORT      Tx       10   164 02fffffd,00fffffd,02acffff,10000000
12:39:04.322  INTR      pstate   10   LR1
12:39:04.323  INTR      pstate   10   LR3
12:39:04.323  INTR      pstate   10   AC
12:39:04.323  PORT      scn     10   11  00000000,00000000,00000002
sw094135:root>
```

See “[Viewing the port log](#)” on page 173 for overview information about a `portlogdump`. Refer to the *HP StorageWorks Fabric OS 4.x command reference guide* for information about decoding a `portlogdump`.

## Identifying media-related issues

This section provides procedures that help pinpoint any media-related issues in the fabric. The tests listed in [Table 32](#) are a combination of *structural* and *functional* tests that can be used to provide an overview of the hardware components and help identify media-related issues.

- *Structural* tests perform basic testing of the switch circuit. If a structural test fails, replace the main board or port card.
- *Functional* tests verify the intended operational behavior of the switch by virtue of running frames through ports or bypass circuitry.

**Table 32** Component test descriptions

Test name	Operands	Checks
<code>crossporttest</code>	<code>[-nframes count]</code> <code>[-lb_mode mode][-spd_mode mode]</code> <code>[-gbic_mode mode] [-norestore mode]</code> <code>[-ports itemlist]</code>	Functional test of port external transmit and receive path.  The <code>crossport</code> is set to loopback using an external cable by default. However, this command can be used to check internal components by setting the <code>lb</code> operand to 5.
<code>fporttest</code>	<code>[-nframes count] [-ports itemlist]</code> <code>[-seed payload_pattern]</code> <code>[-width pattern_width] [-size pattern_size]</code>	Tests component to/from and HBA. Used to test online F_Port devices, N_Port devices, SFPs, and GBICs.
<code>loopporttest</code>	<code>[-nframes count]</code> <code>[-ports itemlist] [-seed payload_pattern]</code> <code>[-width pattern_width]</code>	Only tests components attached to a switch that are on a FC-Al.
<code>spinfab</code>	<code>[nMillionFrames [, ePortBeg [, ePortEnd [, setFail]]]]</code>	Tests components to/from a neighbor switch, such as ISLs, SFPs, and GBICs between switches.

The following procedures are for checking switch-specific components.

To test a port’s external transmit and receive path:

1. Connect to the switch and log in as admin.
2. Connect the port you want to test to any other switch port with the cable you want to test.
3. Issue the `crossporttest` command with the following operands (this is a partial list; refer to the *HP StorageWorks Fabric OS 4.x command reference guide* for additional command information):
  - `[-nframes count]` specifies the number of frames to send.
  - `[-lb_mode mode]` selects the loopback point for the test.
  - `[-spd_mode mode]` selects the speed mode for the test.
  - `[-ports itemlist]` specifies a list of user ports to test.

Example:

```
switch:admin> crossporttest
Running Cross Port Test .... passed.
```

To test a switch's internal components:

1. Connect to the switch and log in as admin.
2. Connect the port you want to test to any other switch port with the cable you want to test.
3. Issue the `crossporttest -lb_mode 5` command.

Where 5 is the operand that causes the test to be run on the internal switch components (this is a partial list—refer to the *Fabric OS Command Reference Manual* for additional command information):

- `[-nframes count]` specifies the number of frames to send.
- `[-lb_mode mode]` selects the loopback point for the test.
- `[-spd_mode mode]` selects the speed mode for the test.
- `[-ports itemlist]` specifies a list of user ports to test.

To test components to and from the HBA:

1. Connect to the switch and log in as admin.
2. Issue the `fPortTest` command (refer to the *HP StorageWorks Fabric OS 4.x command reference guide* for information on the command options).

Example:

```
switchname:admin> fporttest 100,8,0xaa55,2, 512
Will use pattern: aa55 aa55 aa55 aa55 aa55 aa55 ...
Running fPortTest .....
port 8 test passed.
value = 0
```

This example executed the `fPortTest` command 100 times on port 8 with payload pattern 0xaa55, pattern width 2 (meaning word width) and a default payload size of 512 bytes.

See [Table 33](#) for a list of additional tests that can be used to determine the switch components that are not functioning properly. Refer to the *HP StorageWorks Fabric OS 4.x command reference guide* for additional command information.

**Table 33** Switch component tests

Test	Function
portloopbacktest	Functional test of port N to N path.
portregtest	A read and write test of the ASIC SRAMs and registers.
spinsilk	Functional test of internal and external transmit and receive paths at full speed.
sramretentiontest	Verifies that the data written into the miscellaneous SRAMs in the ASIC are retained after a 10-second wait.
crossporttest	Verifies the functional components of the switch.

**Table 33** Switch component tests (continued)

Test	Function
turboramtest	Verifies that the on chip SRAM located in the 2 Gbps ASIC is using the Turbo-Ram BIST circuitry. These same SRAMs are tested by <code>portregtest</code> and <code>sramretentiontest</code> using PCI operations, but for this test the BIST controller is able to perform the SRAM write and read operations at a much faster rate.
statstest	Verifies that the ASIC statistics counter logic.
Related Switch Test Command:	
itemlist	Restricts the items to be tested to a smaller set of parameter values that you pass to the switch.

## Correcting link failures

A link failure occurs when a server or storage is connected to a switch, but the link between the server/storage and the switch does not come up. This prevents the server/storage from communicating through the switch.

If the `switchshow` command and/or the LEDs indicate that the link has not come up properly, use one or more of the following procedures.

To determine if the negotiation was successfully completed:

The port negotiates the link speed with the opposite side. The negotiation usually completes in 1-2 seconds; however, sometimes the speed negotiation fails.



**NOTE:** Skip this procedure if the port speed is set to a static speed through the `portCfgSpeed` command.

1. Issue the `switchshow` command to determine the speed of the device.
2. Issue the `portLogShow` or `portLogDump` command.
3. Check the events area of the output. The first example is 1 Gbps and the second example is 2 Gbps:

```
14:38:51.976  SPEE sn <Port#>  NC  00000001,00000000,00000001
```

```
14:39:39.227  SPEE      sn      <Port#>  NC  00000002,00000000,00000001
```

- The `sn` field indicates a speed negotiation.
- The `NC` field indicates Negotiation Complete.
- The 01 or 02 fields indicate the speed that has been negotiated.

If these fields do not appear, proceed to the [step 4](#).

4. Correct the negotiation by entering the `portCfgSpeed [slotnumber/]portnumber, speed_level` command if the fields in [Table 3](#) do not appear.

To check for a link initialization failure (loop):

1. Verify the port is an L\_Port.
  - a. Issue the `switchShow` command.
  - b. Check the comment field of the output to verify that the switch port indicates an L\_Port. If a loop device is connected to the switch, the switch port must be initialized as an L\_Port.
2. Verify the loop initialization if the port is not an L\_port.
  - a. Issue the `portLogShow` or `portLogDump` command.
  - b. Check the event area for a `loopscn` entry with command code LOOP.

Example:

```
14:35:12.866 tReceive loopscn <Port#> LOOP 10f5cbc0
```

The `loopscn` entry display indicates that the loop initialization is complete.

3. Skip point-to-point initialization.

The switch changes to point-to-point initialization after the Loop Initialization Soft Assigned (LISA) phase of the loop initialization. This behavior sometimes causes trouble with old HBAs. If this is the case, then:

Skip point-to-point initialization by using the `portCfgLport` Command.

To check for a point-to-point initialization failure:

1. Confirm that the port is active.

If a fabric device or another switch is connected to the switch, the switch port must be active.

2. Issue the `portLogShow` or `portLogDump` commands.
3. Verify that the State Change Notification (SCN) code is 1. An SCN of 1 indicates that the port is active.

Example:

```
13:25:12.506 PORT      scn      <Port#>      1
```

4. Skip over the loop initialization phase.

After becoming an active port, the port becomes an F\_Port or an E\_Port depending on the device on the opposite side. If the opposite device is a fabric device, the port becomes an F\_Port. If the opposite device is another switch, the port becomes an E\_Port.

Some fabric devices have problems with loop initialization. If this is evident, issue the following command:

```
portcfgport port #, 1
```

To correct a port that has come up in the wrong mode:

1. Issue the `switchShow` command.

- See the comment fields in Table 34 and follow the suggested actions.

**Table 34** SwitchShow output and suggested action

Output	Suggested action
Disabled	Issue the portEnable command.
Bypassed	Check the output from the portLogShow or portLogDump commands and identify the link initialization stage where the initialization procedure went wrong.
Loopback	Check the output from portLogShow/PortLogDump commands and identify the link initialization stage where the initialization procedure went wrong.
E_Port	If the opposite side is not another switch, the link has come up in a wrong mode. Check the output from the portLogShow/PortLogDump commands and identify the link initialization stage where the initialization procedure went wrong.
F_Port	If the opposite side of the link is a fabric device, the link has come up in a wrong mode. Check the output from portLogShow or PortLogDump commands.
G_Port	The port has not come up as an E_Port or F_Port. Check the output from portLogShow or PortLogDump commands and identify the link initialization stage where the initialization procedure went wrong.
L_Port	If the opposite side is <i>not</i> a loop device, the link has come up in a wrong mode. Check the output from portLogShow or PortLogDump commands and identify the link initialization stage where the initialization procedure went wrong.

## Correcting marginal links

A marginal link involves the connection between the switch and the edge device. Isolating the exact cause of a marginal link involves analyzing and testing many of the components that make up the link (including the switch port, switch SFP, cable, the edge device, and the edge device SFP).

To troubleshoot a marginal link:

- Issue the portErrShow command.

**Example:**

```
switch:admin> porterrshow
      frames   enc   crc   too   too   bad   enc disc link loss loss frjt fbsy
      tx      rx     in    err   shrt  long   eof   out    c3 fail sync sig
sig=====
 0:  22   24   0   0   0   0   0   1.5m   0   7   3   0   0   0
 1:  22   24   0   0   0   0   0   1.2m   0   7   3   0   0   0
 2:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
 3:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
 4: 149m 99m  0   0   0   0   0   448   0   7   6   0   0   0   0
 5: 149m 99m  0   0   0   0   0   395   0   7   6   0   0   0   0
 6: 147m 99m  0   0   0   0   0   706   0   7   6   0   0   0   0
 7: 150m 99m  0   0   0   0   0   160   0   7   5   0   0   0   0
 8:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
 9:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
10:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
11:   0   0   0   0   0   0   0   0     0   0   0   0   2   0   0
12:   0   0   0   0   0   0   0   0     0   0   0   0   2   0   0
13:   0   0   0   0   0   0   0   0     0   0   0   0   2   0   0
14:   0   0   0   0   0   0   0   0     0   0   0   0   2   0   0
15:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
32:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
33:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
34:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
35:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
36:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
37:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
38:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
39:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
40: 99m 146m  0   0   0   0   0   666   0   6   796   7   0   0
41: 99m 149m  0   0   0   0   0   15k   0   2   303   4   0   0
42: 99m 152m  0   0   0   0   0   665   0   2   221   5   0   0
43: 99m 147m  0   0   0   0   0   16k   0   2   144   4   0   0
44:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
45:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
46:   0   0   0   0   0   0   0   0     0   0   0   0   2   0   0
47:   0   0   0   0   0   0   0   0     0   0   0   0   0   0   0
switch:admin>
```

2. Establish if there are a relatively high number of errors (such as CRC errors or ENC\_OUT errors), or if there are a steadily increasing number of errors to confirm a marginal link.
3. If you suspect a marginal link, isolate the areas by moving the suspected marginal port cable to a different port on the switch.  
If the problem stops or goes away, the switch port or the SFP is marginal (proceed to [step 4](#))  
If the problem does *not* stop or go away, see [step 7](#).
4. Replace the SFP on the marginal port.
5. Issue the `portLoopBack` test on the marginal port (refer to the *HP StorageWorks Fabric OS 4.x command reference guide* or see "[Correcting I2C bus errors](#)" on page 190 for additional information).
6. Check the results of the loopback test, and proceed as follows:
  - If the loopback test failed, the port is bad. Replace the port card.
  - If the loopback test did not fail, the SFP was bad.

7. Optionally, to rule out cabling issues:
  - a. Insert a new cable in to the suspected marginal port.
  - b. Issue the `portErrShow` command to determine if a problem still exists.
    - If the `portErrShow` output displays a normal number of generated errors, the issue is solved.
    - If the `portErrShow` output still displays a high number of generated errors, follow the troubleshooting procedures for the Host or Storage device.

---

## Inaccurate information in the system message log

In rare instances, events gathered by the *track change* feature can report inaccurate information to the system message log.

For example, a user enters a correct user name and password, but the login was rejected because the maximum number of users had been reached. However, when looking at the system message log, the login was reported as successful.

If the maximum number of switch users has been reached, the switch performs correctly in that it rejects the login of additional users (even if they enter correct user name and password information).

However, in this limited example, the Track Change feature reports this event inaccurately to the system message log; it appears that the login was successful. This scenario only occurs when the maximum number of users has been reached; otherwise, the login information displayed in the system message log should reflect reality.

For information regarding enabling and disabling Track Changes (TC), see "[Tracking and controlling switch changes](#)" on page 35.

---

## Recognizing the port initialization and FCP auto discovery process

The steps in the port initialization process represent a protocol used to discover the type of connected device and establish the port type. The possible port types are as follows:

- U\_Port      Universal FC port. This port type is the base Fibre Channel port type and all unidentified, or uninitiated ports are listed as U\_Ports.
- FL\_Port     Fabric Loop port. This port connects both public and private loop devices.
- G\_Port      Generic port. This port acts a transition port for non-loop fabric capable devices (E\_Port / F\_Port).
- E\_Port      Expansion port. This port type is assigned to ISL links.
- F\_Port      Fabric port. This port is assigned to fabric capable devices.

The HP FCP auto discovery process enables private storage devices that accept PRLI to communicate in a fabric.

If device probing is enabled, the embedded port PLOGIs and attempts a PRLI into the device to retrieve information to enter into the Name Server. This enables private devices that do not FLOGI but accept PRLI to be entered in the Name Server and receive full fabric citizenship. Private devices that accept PRLI represent a majority of storage targets. Private hosts require the QuickLoop feature, which is not available in Fabric OS v4.0.0 or later.

A fabric-capable device implicitly registers information with Name Server during a FLOGI. These devices typically register information with the Name Server before querying for a device list. The embedded port still PLOGIs and attempts PRLI with these devices.

You can view the Name Server table in Advanced Web Tools by selecting the Name Server button in the Fabric Toolbar. Refer to the *HP StorageWorks Fabric OS 4.x Advanced Web Tools user guide* for more information.



# A Configuring the PID format

This appendix contains the following sections:

- [About PIDs and PID binding](#), page 203
- [Summary of PID formats](#), page 204
- [Impact of changing the fabric PID format](#), page 204
- [Selecting a PID format](#), page 206
- [Evaluating the fabric](#), page 207
- [Planning the update procedure](#), page 209
- [Changing to core PID format](#), page 211
- [Changing to extended edge PID format](#), page 212
- [Performing PID format changes](#), page 213
- [Swapping port area IDs](#), page 217

Port identifiers (called *PIDs*) are used by the routing and zoning services in Fibre Channel fabrics to identify ports in the network. All devices in a fabric must use the same PID format, so when you add new equipment to your SAN, you may need to change the PID format on legacy equipment.

## About PIDs and PID binding

The PID is a 24-bit address built from three 8-bit fields:

- domain
- area\_ID
- AL\_PA

Many scenarios cause a device to receive a new PID; for example, unplugging the device from one port and plugging it into a different port as part of fabric maintenance, or changing the domain ID of a switch, which may be necessary when merging fabrics, or changing compatibility mode settings.

Some device drivers use the PID to map logical disk drives to physical Fibre Channel counterparts. Most drivers can either change PID mappings dynamically (called *dynamic PID binding*) or use the WWN of the Fibre Channel disk for mapping (called *WWN binding*).

Some older device drivers behave as if a PID uniquely identifies a device (they use *static PID binding*). These device drivers should be updated, if possible, to use WWN or dynamic PID binding instead, because static PID binding creates problems in many routine maintenance scenarios. Fortunately, very few device drivers still behave this way. Many current device drivers enable you to select static PID binding as well as WWN binding. You should only select static binding if there is a compelling reason, and only after you have evaluated the impact of doing so.

# Summary of PID formats

HP StorageWorks switches employ these types of PID formats:

- VC encoded

This is the format defined by the Fibre Channel Storage Switch 8 and Fibre Channel Storage Switch 16. Connections to these switches are not supported in Fabric OS v4.0.0 and later.

- native

Introduced with the HP StorageWorks 1 GB switches, this format supports up to 16 ports per switch.

- core

The default for HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, SAN Switch 4/32, Core Switch 2/64, and SAN Director 2/128, this is the recommended format for HP StorageWorks switches and fabrics. It uses the entire 8-bit address space and directly uses the port number as the area\_ID. It supports up to 256 ports per switch.

- extended edge

This format generates the same PID for a port on switches with 16 ports or less as would native PID format, but it also supports up to 128 ports per domain. It should be used only in cases where you cannot upgrade devices to dynamic PID binding and you absolutely cannot reboot your servers.

Extended edge PID is supported in Fabric OS v2.6.2 and later, v3.1.2 and later, and v4.2.0 and later.

In addition to the PID formats list here, Interoperability mode supports additional PID formats that are not discussed in this guide.

## Impact of changing the fabric PID format

If your fabric contains switches that use Native PID, HP recommends that you change the format to Core PID before you add the new, higher port count switches and directors. Also, HP recommends that you use Core PID when upgrading the Fabric OS version on HP StorageWorks 1 GB and 2 GB switches.

Depending on your situation, the PID change may or may not entail fabric downtime:

- If you are running dual-fabrics with multipathing software, you can update one fabric at a time without disrupting traffic. Move all traffic onto one fabric in the SAN and update the other fabric. Then move the traffic onto the updated fabric, and update the final fabric.
- Without dual-fabrics, stopping traffic is highly recommended. This is the case for many routine maintenance situations, so dual-fabrics are always recommended for uptime-sensitive environments. If your fabric contains devices that employ static PID binding, or you do not have dual-fabrics, you must schedule downtime for the SAN to change the PID format.

The following sections describe various impacts of PID format changes in greater detail.

## Host reboots

In some Fibre Channel SAN environments, storage devices and host servers are bound to the host operating system by their PIDs (called their *Fibre Channel addresses*). In these environments, the hosts and target HBAs in a SAN need to know the full 24-bit PIDs of the hosts and targets they are communicating with, but they do not care how the PIDs are determined. But, if a storage device PID is changed, the host must reestablish a new binding, which requires the host to be rebooted.

With the introduction of the HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, SAN Switch 4/32, Core Switch 2/64, and SAN Director 2/128, the Native PID format used in earlier switches was supplemented with the Core PID format, which is capable of addressing higher port counts. Changing from Native PID format to Core PID format changes the PID, which requires hosts that use port binding to be rebooted.

## Static PID mapping errors

If you can avoid using drivers that employ static PID binding, you should do so.

With the WWN or dynamic PID binding most typically used with drivers, changing the device's PID does not affect the PID mapping. However, before updating the PID format, it is necessary to determine whether any devices in the SAN use static PID binding.

For those few drivers that do use static PID binding, changing the PID format breaks the mapping, which must be fixed either by rebooting the host or by using a manual update procedure on the host.

To correct mapping errors caused by static PID binding, refer to the following sections:

- See "[Evaluating the fabric](#)" on page 207 for details on finding devices that use static PID binding. Then see "[Online update](#)" on page 209 or "[Offline update](#)" on page 210 for recommendations.
- See "[Performing PID format changes](#)" on page 213 for instructions.

## Changes to configuration data

[Table 35](#) lists various combinations of before-and-after PID formats, and indicates whether the configuration is affected.



**NOTE:** After changing the fabric PID format, if the change invalidates the configuration data (see [Table 35](#) to determine this), do not download old (pre-PID format change) configuration files to any switch on the fabric.

**Table 35** Effects of PID format changes on configurations

PID format before change	PID format after change	Configuration effect?
Native	Extended Edge	No impact
Extended Edge	Native	No impact

**Table 35** Effects of PID format changes on configurations (continued)

PID format before change	PID format after change	Configuration effect?
Native	Core	You must: <ul style="list-style-type: none"><li>• Reenable zoning, if there is an active zone set and it uses port zones.</li></ul>
Core	Native	
Extended Edge	Core	
Core	Extended Edge	<ul style="list-style-type: none"><li>• If Destination ID (DID) binding is used, reconfigure persistent binding, and reconfigure DID list for performance monitor.</li></ul>

After changing the fabric PID format and verifying correct fabric operation, resave configuration data by running the `configUpload` command.

Before downgrading firmware, change the PID back to supported PIDs such as Core PID. If the database is automatically converted, save the converted database, and then download the older OS.

## Selecting a PID format

All switches in a fabric must use the same PID format, so if you add a switch that uses a different PID format to a fabric, the switch segments from the fabric. The format you select for your fabric depends on the mix of switches in the fabric, and to an extent on the specific releases of Fabric OS in use (for example, Extended Edge PID format is only available in Fabric OS v2.6.2 and later, Fabric OS v3.1.2 and later, and Fabric OS v4.2.0 and later).

If you are building a new fabric with switches running various Fabric OS versions, use Core PID format to simplify port-to-area\_ID mapping.

**Table 36** shows various combinations of existing fabrics, new switches added to those fabrics, and the recommended PID format for that combination. The criteria for the recommendations are first to eliminate host reboots, and second to minimize the need for a host reboot in the future.

**Table 36** PID format recommendations for adding new switches

Existing Fabric OS versions; PID format	Switch to be added	Recommendations (in order of preference)
v2.0.0 and later/v3.1.2 and later; Native PID	v2.0.0 and later/v3.1.2 and later	<ol style="list-style-type: none"> <li>1. Use Native PID format for new switch Host reboot is not required.</li> <li>2. Convert existing fabric to Core PID format, upgrading the version of Fabric OS, if necessary. Set Core PID format for new switch. Host reboot is required.</li> <li>3. If devices are bound statically and it is not possible to reboot, convert existing fabric to Extended Edge PID format, upgrading the version of Fabric OS, if necessary. Use Extended Edge PID format for new switch Host reboot is not required.</li> </ol>
	v4.2.0 and later	<ol style="list-style-type: none"> <li>1. Convert existing fabric to Core PID format, upgrading the version of Fabric OS, if necessary. Set Core PID format for new switch. Host reboot is required.</li> <li>2. If devices are bound statically and it is not possible to reboot, convert existing fabric to Extended Edge PID format, upgrading the version of Fabric OS, if necessary. Use Extended Edge PID format for new switch Host reboot is not required.</li> </ol>
v2.0.0 and later/v3.1.2 and later/v4.2.0 and later; Core PID	v2.0.0 and later/v3.1.2 and later/v4.2.0 and later	Use Core PID for new switch Host reboot is not required.

## Evaluating the fabric

If there is the possibility that your fabric contains host devices with static PID bindings, you should evaluate the fabric to:

- Find any devices that bind to PIDs
- Determine how each device driver responds to the PID format change
- Determine how any multipathing software responds to a fabric service interruption

If current details about the SAN are already available, it may be possible to skip the Data Collection step. If not, it is necessary to collect information about each device in the SAN. Any type of device may be able to bind by PID; each device should be evaluated before attempting an online update. This information has broad applicability, because PID-bound devices are not able to seamlessly perform in many routine maintenance or failure scenarios.

**1.** Collect device, software, hardware, and configuration data.

The following is a non-comprehensive list of information to collect:

- HBA driver versions
- Fabric OS versions
- RAID array microcode versions
- SCSI bridge code versions
- JBOD drive firmware versions
- Multipathing software versions
- HBA time-out values
- Multipathing software time-out values
- Kernel time-out values
- Configuration of switch

**2.** Make a list of manually configurable PID drivers.

Some device drivers do not automatically bind by PID, but allow the operator to manually create a PID binding. For example, persistent binding of PIDs to logical drives may be done in many HBA drivers. Make a list of all devices that are configured this way. If manual PID binding is in use, consider changing to WWN binding.

The following are some of the device types that may be manually configured to bind by PID:

- HBA drivers (persistent binding)
- RAID arrays (LUN access control)
- SCSI bridges (LUN mapping)

**3.** Analyze data.

After you have determined the code versions of each device on the fabric, they must be evaluated to find out if any automatically bind by PID. It may be easiest to work with the support providers of these devices to get this information. If this is not possible, you may need to perform empirical testing.

Binding by PID can create management difficulties in a number of scenarios. HP recommends that you not use drivers that bind by PID. If the current drivers do bind by PID, upgrade to WWN-binding drivers if possible.

The drivers shipping by default with HP/UX and AIX at the time of this writing still bind by PID, and so detailed procedures are provided for these operating systems in this chapter. Similar procedures can be developed for other operating systems that run HBA drivers that bind by PID.

There is no inherent PID binding problem with either AIX or HP/UX. It is the HBA drivers shipping with these operating systems that bind by PID. Both operating systems are expected to release HBA drivers that bind by WWN, and these drivers may already be available through some support channels. Work with the appropriate support provider to find out about driver availability.

It is also important to understand how multipathing software reacts when one of the two fabrics is taken offline. If the time-outs are set correctly, the failover between fabrics should be transparent to the users.

You should use the multipathing software to manually fail a path before starting maintenance on that fabric.

#### 4. Perform empirical testing.

Empirical testing may be required for some devices, to determine whether they bind by PID. If you are not sure about a device, work with the support provider to create a test environment.

Create as close a match as practical between the test environment and the production environment, and perform an update using the procedure in "[Online update](#)" on page 209.

Devices that bind by PID are unable to adapt to the new format, and one of three approaches must be taken with them:

- A plan can be created for working around the device driver's limitations in such a way as to allow an online update. See the Detailed Procedures section for examples of how this could be done.
- The device can be upgraded to drivers that do not bind by PID.
- Downtime can be scheduled to reset the device during the core PID update process, which generally allows the mapping to be rebuilt.

If either of the first two options are used, the procedures should again be validated in the test environment.

Determine the behavior of multipathing software, including but not limited to:

- HBA time-out values
- Multipathing software time-out values
- Kernel time-out values

---

## Planning the update procedure

Whether it is best to perform an offline or online update depends on the uptime requirements of the site.

- An offline update that all devices attached to the fabric be offline.
- With careful planning, it should be safe to update the core PID format parameter in a live, production environment. This requires dual fabrics with multipathing software. Avoid running backups during the update process, as tape drives tend to be very sensitive to I/O interruption. The online update process is only intended for use only in uptime-critical dual-fabric environments, with multipathing software (high-uptime environments should always use a redundant fabric SAN architecture). Schedule a time for the update when the least critical traffic is running.

All switches running any version of Fabric OS 3.1.2 and later or 4.2.0 and later are shipped with the Core Switch PID Format enabled, so it is not necessary to perform the PID format change on these switches.

Migrating from manual PID binding (such as persistent binding on an HBA) to manual WWN binding and upgrading drivers to versions that do not bind by PID can often be done before setting the core PID format. This reduces the number of variables in the update process.

## Online update

The following steps are intended to provide SAN administrators a starting point for creating site-specific procedures.

1. Back up all data and verify backups.
2. Verify that the multipathing software can automatically switchover between fabrics seamlessly. If there is doubt, use the software's administrative tools to manually disassociate or mark offline all storage devices on the first fabric to be updated.

3. Verify that I/O continues over the other fabric.
4. Disable all switches in the fabric to be updated, one switch at a time, and verify that I/O continues over the other fabric after each switch disable.
5. Change the PID format on each switch in the fabric.
6. Reenable the switches in the updated fabric one at a time. In a core/edge network, enable the core switches first.
7. After the fabric has reconverged, use the `cfgenable` command to update zoning.
8. Update their bindings for any devices manually bound by PID. This may involve changing them to the new PIDs, or preferably changing to WWN binding.

For any devices automatically bound by PID, two options exist:

- a. Execute a custom procedure to rebuild its device tree online. Examples are provided in the “[Performing PID format changes](#)” on page 213 section of this chapter.
- b. Reboot the device to rebuild the device tree. Some operating systems require a special command to do this, for example `boot -r` in Solaris.

9. For devices that do not bind by PID or have had their PID binding updated, mark online or reassociate the disk devices with the multipathing software and resume I/O over the updated fabric.
10. Repeat with the other fabric(s).

## Offline update

The following steps are intended to provide SAN administrators a starting point for creating site-specific procedures.

1. Schedule an outage for all devices attached to the fabric.
2. Back up all data and verify backups.
3. Shut down all hosts and storage devices attached to the fabric.
4. Disable all switches in the fabric.
5. Change the PID format on each switch in the fabric.
6. Reenable the switches in the updated fabric one at a time. In a core/edge network, enable the core switches first.
7. After the fabric has reconverged, use the `cfgenable` command to update zoning.
8. Bring the devices online in the order appropriate to the SAN. This usually involves starting up the storage arrays first, and the hosts last.
9. For any devices manually bound by PID, bring the device back online, but do not start applications. Update their bindings and reboot again if necessary. This may involve changing them to the new PIDs, or may (preferably) involve changing to WWN binding.
10. For any devices automatically bound by PID, reboot the device to rebuild the device tree (some operating systems require a special command to do this, such as `boot -r` in Solaris).
11. For devices that do not bind by PID or have had their PID binding updated, bring them back up and resume I/O.
12. Verify that all I/O has resumed correctly.

## Hybrid update

It is possible to combine the online and offline methods for fabrics where only a few devices bind by PID. Because any hybrid procedure is extremely customized, it is necessary to work closely with the SAN service provider in these cases.

# Changing to core PID format

In Fabric OS release v4.2.0 and later, Native PID format is not supported; the default format is the Core PID format.

In Fabric OS v3.1.2 and later, Core PID format is the default configuration.

In Fabric OS v2.0.0 and later, Native PID format is the default configuration.

Although the PID format is listed in the configuration file, do not edit the file to change the setting there. Instead, use the CLI `configure` command. When you use the `configure` command, switch databases that contain PID-sensitive information are automatically updated. If you change the setting in the configuration file and then download the edited file, the PID format is changed, but the database entries are not, and so they are incorrect.

[Table 37](#) maps the PID format names to the names used in the management interfaces.

**Table 37** PID format and management interface names

PID format name	Management interface name
native PID	switch PID address mode 0
core PID	switch PID address mode 1
extended edge PID	switch PID address mode 2

Before changing the PID format, determine if host reboots are necessary. The section "[Host reboots](#)" on page 205 summarizes the situations that may require a reboot.

Example:

```
switch:admin> switchdisable
switch:admin> configure
Configure...

Fabric parameters (yes, y, no, n): [no] y

Domain: (1..239) [1]
BB credit: (1..27) [16]
R_A_TOV: (4000..120000) [10000]
E_D_TOV: (1000..5000) [2000]
WAN_TOV: (1000..120000) [0]
Data field size: (256..2112) [2112]
Sequence Level Switching: (0..1) [0]
Disable Device Probing: (0..1) [0]
Suppress Class F Traffic: (0..1) [0]
SYNC IO mode: (0..1) [0]
Switch PID Address Mode: (0..2) [1] < Set mode number here.
Per-frame Route Priority: (0..1) [0]
Long Distance Fabric: (0..1) [0]
```

# Changing to extended edge PID format

In rare cases, you may be affected by the presence in the fabric of drivers that rely on static binding to the dynamically assigned PID; for example, you may be installing a switch running Fabric OS v4.2.0 into a fabric consisting solely of Fabric OS v2.0.0 and later or v3.1.2 and later switches. In these cases, if you absolutely cannot reboot the affected servers when you upgrade your switches, you can choose Extended Edge PID format. It uses the same PID mapping for the first 16 ports and can support switches and directors with higher port counts. However, because Extended Edge format only supports 128 ports per domain, its use can lead to port addressing issues in directors.

Use the following procedure only if your fabric contains devices that are bound statically and you cannot reboot the host.

1. Determine if the current switch firmware versions meet the minimum supported version levels.

[Table 38](#) lists the earliest Fabric OS version levels that support Extended Edge PID format. Use this table to determine if you need to upgrade the firmware in the switches in your fabric before you change the PID format.

**Table 38** Earliest Fabric OS versions for extended edge PID format

1 GB switches	2 GB switches: SAN Switches 2/8-EL and 2/16	2 GB switches: SAN Switches 2/8V, 2/16V, 2/32, Core Switch 2/64, and SAN Director 2/128
2.6.2	3.1.2	4.2.0

2. Update switch firmware as necessary.
  - a. Issue the `fabricshow` command to verify the total number of switches in the fabric.
  - b. Download the correct firmware version to each switch as necessary.
  - c. Reboot all switches.
  - d. Verify that the switches form a single fabric and that all domain IDs do not change after forming the fabric.
  - e. Verify that the number of switches is the same.
3. Disable the switch by issuing the `switchdisable` command.
4. Change the switch configuration in the fabric to Extended Edge PID format.
  - a. Configure Extended Edge PID (Format 2) on each switch. (See [Figure 8](#) for a sample configure command on a HP StorageWorks switch running Fabric OS v3.1.2 and later and see [Figure 9](#) for a sample configure command on an HP StorageWorks switch running Fabric OS 4.2.0 and later.)
  - b. Issue the `switchenable` command on all switches.
  - c. Verify that all the switches form a fabric.
  - d. Issue the `switchshow` command to verify the interswitch links (ISLs) are correct and the device links are correct.
  - e. Issue the `fabricshow` command to verify that the number switches are the same as those when starting this procedure.
  - f. Issue the `nsallshow` command to verify the total number of devices is the same as those when starting this procedure.
5. For dual fabrics, repeat step 1 through step 4 for the other fabric.

```
Configure...

Fabric parameters (yes, y, no, n): [no] yes

Domain: (1..239) [217]
BB credit: (1..27) [16]
R_A_TOV: (4000..120000) [10000]
E_D_TOV: (1000..5000) [2000]
Data field size: (256..2112) [2112]
Sequence Level Switching: (0..1) [0]
Disable Device Probing: (0..1) [0]
Suppress Class F Traffic: (0..1) [0]
SYNC IO mode: (0..10 [0]
Switch PID Format : (0..2) [0] 2
Per-frame Route Priority: (0..1) [0]
Long Distance Fabric: (0..1) [0]

Virtual Channel parameters (yes, y, no, n): [no] ^D
Committing configuration...done.
0x102fd500 (tshell): Apr 15 16:53:31
WARNING CONFIG-PIDCHANGE_DISPLACE, 3, Switch PID format changed to
Extended Edge PID Format
```

**Figure 8** Configure command on a switch running Fabric OS 3.1.2

```
Configure...

Fabric parameters (yes, y, no, n): [no] yes

Domain: (1..239) [112]
R_A_TOV: (4000..120000) [10000]
E_D_TOV: (1000..5000) [2000]
Data field size: (256..2112) [2112]
Sequence Level Switching: (0..1) [0]
Disable Device Probing: (0..1) [0]
Suppress Class F Traffic: (0..1) [0]
Switch PID Format: (1..2) [1] 2
Per-frame Route Priority: (0..1) [0]
Long Distance Fabric: (0..1) [0]
BB credit: (1..16) [16]
```

**Figure 9** Configure command on a switch running Fabric OS 4.2.0 and later

## Performing PID format changes

There are several routine maintenance procedures which may result in a device receiving a new PID. Examples include, but are not limited to:

- Changing compatibility mode settings
- Changing switch domain IDs
- Merging fabrics

- Relocating devices to new ports or new switches (that is, for Add, Move, Change type operations)
- Updating the core PID format
- Using hot spare switch ports to deal with failures

In every case where devices employ static PID binding, any such procedure becomes difficult or impossible to execute without downtime.

In some cases, device drivers allow you to specify static PID binding. In these cases, such devices must be identified and their PID binding should be changed to WWN binding.

The following sections contain a basic procedure that summarizes the steps necessary to perform PID format changes without disrupting the fabric, and special procedures for HP/UX and AIX.

## Basic procedure

This process should be executed as part of the overall online or offline update process. However, it can be implemented in a stand-alone manner on a non-production fabric, or a switch that has not yet joined a fabric.

1. Ensure that all switches in the fabric are running Fabric OS versions that support the addressing mode. HP recommends that you use v2.6.2 for HP StorageWorks 1 GB switches; v3.1.2 for SAN Switch 2/8-EL and SAN Switch 2/16; v4.2.0 for HP StorageWorks 2GB switches, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64 and SAN Director 2/128.



**NOTE:** All switches running any version of Fabric OS 4.0.0 and later are shipped with the Core Switch PID Format enabled, so it is not necessary to perform the PID format change on these switches.

2. Telnet into one of the switches in the fabric.
3. Disable the switch by issuing the `switchdisable` command.
4. Issue the `configure` command (the `configure` prompts display sequentially).
5. Enter `y` after the `Fabric parameters` prompt.
6. Enter `1` at the `Core Switch PID Format` prompt.
7. Complete the remaining prompts or press **Ctrl+d** to accept the remaining settings without completing all the prompts.
8. Repeat step 2 through step 7 for the remaining switches in the fabric.
9. Reenable the switch by entering the `switchenable` command.

Example:

```
switch:admin> switchdisable
switch:admin> configure
Configure...
Fabric parameters (yes, y, no, n): [no] yes
Domain: (1..239) [1]
R_A_TOV: (4000..120000) [10000]
E_D_TOV: (1000..5000) [2000]
Data field size: (256..2112) [2112]
Sequence Level Switching: (0..1) [0]
Disable Device Probing: (0..1) [0]
Suppress Class F Traffic: (0..1) [0]
SYNC IO mode: (0..1) [0]
Core Switch PID Format: (0..2) [0] 1
Per-frame Route Priority: (0..1) [0]
Long Distance Fabric: (0..1) [0]
BB credit: (1..27) [16]
```

**10.** After all switches are updated to use the new PID format and reenabled, verify that the fabric has fully reconverged (each switch sees the other switches).

**11.** Issue the command `cfgenable [active_zoning_config]` on one of the switches in the fabric to update zoning to use the new PID form.

This does not change the definition of zones in the fabric, but merely causes the lowest level tables in the zoning database to be updated with the new PID format setting. It is only necessary to do this once per fabric; the zoning update automatically propagates to all switches.

At this point, all switches in the fabric are operating in the new addressing mode.

## HP-UX procedure

This procedure is not intended to be comprehensive. It provides a starting point from which a SAN administrator could develop a site-specific procedure for a device that binds automatically by PID, and cannot be rebooted due to uptime requirements.

**1.** Backup all data. Verify backups.

**2.** If you are not using multipathing software, stop all I/O going to all volumes connected through the switch/fabric to be updated.

**3.** If you are not using multipathing software, unmount the volumes from their mount points using `umount`. The format is `umount mount_point`. For example:

```
umount /mnt/jbod
```

**4.** If you are using multipathing software, use that software to remove one fabric's devices from its configuration.

**5.** Deactivate the appropriate volume groups using `vgchange`. The format is `vgchange -a n path_to_volume_group`. For example:

```
vgchange -a n /dev/jbod
```

**6.** Make a backup copy of the volume group directory using `tar` from within `/dev`. For example:

```
tar -cf /tmp/jbod.tar jbod
```

**7.** Export the volume group using `vgexport`. The format is `vgexport -m mapfile path_to_volume_group`. For example:

```
vgexport -m /tmp/jbod_map /dev/jbod
```

**8.** Connect to each switch in the fabric

**9.** Issue the `switchDisable` command.

**10.** Issue the `configure` command and change the Core Switch PID Format to 1.

**11.** Issue the command `cfgEnable [effective_zone_configuration]`. For example:

```
cfgEnable my_zones
```

**12.** Issue the `switchEnable` command. Enable the core switches first, then the edges.

**13.** Clean the `lvmtab` file by using the command `vgscan`.

**14.** Change to `/dev` and untar the file that was tared in step 4. For example:

```
tar -xf /tmp/jbod.tar
```

**15.** Import the volume groups using `vgimport`. The format is

```
vgimport -m mapfile path_to_volume_group physical_volume_path. For example:
```

```
vgimport -m /tmp/jbod_map /dev/jbod /dev/dsk/c64t8d0 /dev/dsk/c64t9d0
```

**16.** Activate the volume groups using `vgchange`. The format is

```
vgchange -a y path_to_volume_group. For example:
```

```
vgexport -a y /dev/jbod
```

**17.** If you are not using multipathing software, mount all devices again and restart I/O. For example:

```
mount /mnt/jbod
```

**18.** If you are using multipathing software, reenable the affected path. The preceding steps do not clean up the results from ioscan. When viewing the output of ioscan, notice that the original entry is still there, but now has a status of NO\_HW.

#	ioscan -func disk	Class	I	H/W Path	Driver	S/W State	H/W Type	Description
		disk	0	0/0/1/1.2.0	adisk	CLAIMED	DEVICE	SEAGATE ST39204LC /dev/dsk/clt2d0 /dev/rdsk/clt2d0
		disk	1	0/0/2/1.2.0	adisk	CLAIMED	DEVICE	HP DVD-ROM 304 /dev/dsk/c3t2d0 /dev/rdsk/c3t2d0
		disk	319	0/4/0/0.1.2.255.14.8.0	adisk	CLAIMED	DEVICE	SEAGATE ST336605FC /dev/dsk/c64t8d0 /dev/rdsk/c64t8d0
		disk	320	0/4/0/0.1.18.255.14.8.0	adisk	NO_HW	DEVICE	SEAGATE ST336605FC /dev/dsk/c65t8d0 /dev/rdsk/c65t8d0

**19.** To remove the original (outdated) entry, use the rmsf (remove special file) command. The format for this command is rmsf -a -v *path\_to\_device*. For example:

```
rmsf -a -v /dev/dsk/c65t8d0
```

**20.** Validate that the entry has been removed by issuing the ioscan -func disk command. In the following example, the NO\_HW entry is no longer listed:

#	ioscan -func disk	Class	I	H/W Path	Driver	S/W State	H/W Type	Description
		disk	0	0/0/1/1.2.0	adisk	CLAIMED	DEVICE	SEAGATE ST39204LC /dev/dsk/clt2d0 /dev/rdsk/clt2d0
		disk	1	0/0/2/1.2.0	adisk	CLAIMED	DEVICE	HP DVD-ROM 304 /dev/dsk/c3t2d0 /dev/rdsk/c3t2d0
		disk	319	0/4/0/0.1.2.255.14.8.0	adisk	CLAIMED	DEVICE	SEAGATE ST336605FC /dev/dsk/c64t8d0 /dev/rdsk/c64t8d0

**21.** Repeat for all fabrics.

**22.** Issue the switchEnable command. Enable the core switches first, then the edges.

## AIX procedure

This procedure is not intended to be comprehensive. It provides a starting point from which a SAN administrator can develop a site-specific procedure for a device that binds automatically by PID, and cannot be rebooted due to uptime requirements.

1. Backup all data. Verify backups.
2. If you are not using multipathing software, stop all I/O going to all volumes connected through the switch or fabric to be updated.
3. If you are not using multipathing software, varyoff the volume groups. The command format is varyoffvg *volume\_group\_name*. For example:  

```
varyoffvg datavg
```
4. If you are not using multipathing software, unmount the volumes from their mount points using umount. The command format is umount *mount\_point*. For example:  

```
umount /mnt/jbod
```
5. If you are using multipathing software, use that software to remove one fabric's devices from its configuration.

6. Remove the device entries for the fabric you are migrating. For example, if the HBA for that fabric is fcs0, issue the command:

```
rmdev -Rdl fcs0
```

7. Connect to each switch in the fabric.
8. Issue the `switchdisable` command.

9. Issue the `configure` command and change the Core Switch PID Format to 1.

10. Issue the `configenable [effective_zone_configuration]` command. For example:

```
configenable my_config
```

11. Issue the `switchenable` command. Enable the core switches first, then the edges.

12. Rebuild the device entries for the affected fabric using the `cfgmgr` command. For example:

```
cfgmgr -v
```

This command may take several minutes to complete.

13. If you are not using multipathing software, vary on the disk volume groups. The format is `varyonvg volume_group_name`. For example:

```
varyonvg datavg
```

14. If you are not using multipathing software, mount all devices again and restart I/O. For example:

```
mount /mnt/jbod
```

15. If you are using multipathing software, reenable the affected path.

16. Repeat for all fabrics.

---

## Swapping port area IDs

If a device that uses port binding is connected to a port that fails, you can use port swapping to make another physical port use the same PID as the failed port. The device can then be plugged into the new port without the need to reboot the device.

Use the following procedure to swap the port area IDs of two physical switch ports. In order to swap port area IDs, the port swap feature must be enabled, and both switch ports must be disabled. The swapped area IDs for the two ports remain persistent across reboots, power cycles, and failovers.

### Swap area IDs for a pair of switch ports as follows:

1. Connect to the switch and log in as admin.
2. Enable the port swap feature:

```
portswapenable
```

3. For the HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32: issue the following commands:

```
portdisable port1
```

```
portdisable port2
```

For the Core Switch 2/64 and SAN Director 2/128: issue the following commands:

```
portdisable slot/port1
```

```
portdisable slot/port2
```

4. For the HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32: issue the following command:

```
portswap port1 port2
```

For the Core Switch 2/64 and SAN Director 2/128: issue the following command:

```
portswap slot1/port1 slot2/port2
```

5. Verify that the port area IDs have been swapped:

```
portswapshow
```

A table is shows the physical port numbers and the logical area IDs for any swapped ports.

6. Disable the port swap feature:

```
portswapdisable
```

# B Configuring interoperability mode

This appendix contains the following sections:

- [Vendor switch requirements](#), page 219
- [HP StorageWorks switch requirements](#), page 219
- [Supported HP StorageWorks features](#), page 220
- [Unsupported HP StorageWorks features](#), page 220
- [Configuration recommendations](#), page 221
- [Configuration restrictions](#), page 221
- [Enabling and disabling interoperability mode](#), page 222

This appendix provides information on setting up a heterogeneous fabric that includes HP StorageWorks switches and switches from other manufacturers.

The *interoperability* mode enables HP StorageWorks switches and others to exchange interoperability parameters, allowing their fabrics to merge into one fabric with one principal switch and unique domain IDs.

The `interopmode` command must be executed on all HP StorageWorks switches in the fabric. Each HP StorageWorks switch must be rebooted after changing interoperability mode.

Switches from other manufacturers may also require the execution of one or more commands to set up interoperability mode. Refer to their documentation for details.

In a heterogeneous fabric, each HP StorageWorks switch must have interoperability mode enabled (see “[Enabling and disabling interoperability mode](#)” on page 222).

To provide maximum compatibility between switches, several features are not available in heterogeneous fabrics (see “[Supported HP StorageWorks features](#)” on page 220 and “[Unsupported HP StorageWorks features](#)” on page 220).

## Vendor switch requirements

To determine whether or not a mixed-vendor SAN is supported, work with the switch provider to determine if your SAN design is valid. Important variables that determine the supportability of a particular mixed vendor SAN include the number of switches, version of Fabric OS, the topology, number of ISLs, number of connected devices, and hop count. For interoperability rules, refer to *HP StorageWorks SAN design reference guide*.

## HP StorageWorks switch requirements

The following are the HP StorageWorks software requirements:

- 1 GB switches must be running Fabric OS v2.6.0 or later.
- 2 GB switches must be running Fabric OS v3.1.0 or later.

- SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128 must be running 4.2.0 or later.
- SAN Switch 4/32 must be running Fabric OS v4.4.0 or later.
- A zoning license and a fabric license must be installed on each HP StorageWorks switch.

---

## Supported HP StorageWorks features

The following features are supported on HP StorageWorks switches in interoperability mode:

- HP Fabric Watch
- HP Fabric Access API functions

Accessible from HP StorageWorks switches only, but switch information for non-HP StorageWorks switches is reported. The object information and zoning actions are configurable from the API.

- HP translative mode

Registers private storage target devices into the fabric, it can be used in a heterogeneous fabric if the devices are connected directly to HP StorageWorks switches. The devices are accessible from any port on the fabric.

---

## Unsupported HP StorageWorks features

In a heterogeneous fabric, the following HP optional features are not supported and cannot be installed on any HP StorageWorks switch in the fabric:

- Extended Edge PID format
- QuickLoop and QuickLoop Zoning
- Secure Fabric OS
- Timer Server function
- Open E\_Port
- Broadcast Zoning
- Management Server Service and FDMI
- QuickLoop Fabric Assist
- Remote Switch
- Extended Fabrics
- Trunking
- Alias Server
- Platform Service
- Virtual Channels
- FC-IP

# Configuration recommendations

The following is recommended when configuring an interoperable fabric:

- Avoid domain ID conflicts before fabric reconfiguration. Every switch in the fabric must have a unique domain ID.
- When you are configuring multiple switches, you should wait for a fabric reconfiguration after adding or removing each switch.

# Configuration restrictions

In interoperable fabrics, the following restrictions apply:

- Do not use Extended Edge PID mode.
- There is an architecture maximum of 31 switches.
- Domain IDs must be in the 97 to 127 value range for successful connection to McData switches. The firmware automatically assigns a valid domain ID, if necessary, when the `interopmode` command is enabled on the switch.
- The `fabricshow` command only shows the WWN and domain ID for McData switches. No IP address or switch name information is provided. HP StorageWorks switches show all parameters.
- If you are managing zoning from HP StorageWorks switches, then all HP StorageWorks switches must have at least one direct connection to another HP StorageWorks switch. For example, you cannot have a McData switch between two HP StorageWorks switches if you are managing zoning from the HP StorageWorks switches.
- LC IBM GBICs are not supported if they are connected to a McData ISL.
- When an HP StorageWorks switch gets a new domain ID assigned through a fabric reconfiguration, the new domain ID is written to nonvolatile memory and the old domain ID value is overwritten. When a McData switch gets a new domain ID assigned through a fabric reconfiguration, it keeps the original domain ID in nonvolatile memory.

In this scenario, when the domain ID of both a McData switch and an HP StorageWorks switch are changed via fabric reconfiguration, on the next and subsequent fabric reconfigurations, the HP StorageWorks switch attempts to use the new ID (from the nonvolatile memory) while McData attempts to use its old ID (from the nonvolatile memory). This situation may cause a domain ID overlap to occur during multiple fabric reconfigurations. Domain ID overlap is not supported for HP/McData interoperability.

- Between HP StorageWorks switches, you can connect more than one ISL when in interoperability mode.

# Zoning restrictions

The following restrictions apply to zoning in interoperable fabrics:

- When interoperability mode is in effect, the space available for the zoning database is about half the usual size. The maximum zoning database size in interoperability mode is 1:1.8 of the native mode zoning database size
- Only zoning by port WWN is allowed; you must use the port WWN of the device, such as 10:00:00:00:c9:28:c7:c6.
- Zone members specified by node WWN are ignored.

- Zone configurations that use either physical port numbers or port IDs are not supported in interoperability mode. Zoning using port numbers uses the actual physical port numbers on the switch; for example slot 1, port 5.
- When a zoning configuration is not in effect, by default all ports are isolated and traffic is not permitted. This is unlike HP StorageWorks behavior where *Interoperability* mode is off (and all data traffic is enabled).
- HP StorageWorks SAN Switch 2/8-EL, SAN Switch 2/8V, SAN Switch 2/16, SAN Switch 2/16V, SAN Switch 2/32, SAN Switch 4/32, Core Switch 2/64, and SAN Director 2/128 provide hardware enforcement of the port WWN zones only for devices attached to their ports. Devices attached to end-ports on non-HP StorageWorks switches or HP StorageWorks 1 GB switches are enforced by Name Server (soft) zoning only.
- Advanced Web Tools can be used for zone configuration as long as HP StorageWorks switches are connected directly to each other. If Advanced Web Tools is used to set up zoning, then it must be used as the only zone management method.
- HP StorageWorks switches connected to a McData switch receive the effective configuration when a zone merge occurs. (McData only has an effective zone configuration and discards the defined zone configuration when it sends merge information to the HP StorageWorks switch.) However, a zone update sends both the defined and the effective configuration to all switches.
- When an HP StorageWorks switch or director is reconfiguring, wait until the fabric routes are completely set up before entering zoning commands that must propagate to other switches. Use the `fabricshow` command to verify that all fabric routes are set up and all switch IP addresses and names are present. (The `fabricshow` command only shows the WWN and domain ID for switches from other manufacturers.)
- The maximum number of items that can be stored in the zoning configuration database depends on the switches in the fabric, whether or not interoperability mode is enabled, and the number of bytes required for each item.

You can use the `cfgsize` command to check both the maximum available size and the currently saved size. If you believe you are approaching the maximum, you can save a partially completed zoning configuration and use the `cfgSize` command to determine the remaining space.

## Zone name restrictions

The name field must contain the ASCII characters that actually specify the name, not including any required fill bytes. Names must follow these rules:

- Length must be between 1 and 64 characters.
- All characters must be 7-bit ASCII.
- The first character must be a letter, which can be either upper case (A-Z) or lower case (a-z).
- Any character other than the first character must be lower case (a-z), upper case (A-Z), a single-digit number (0-9), dash (-), or underscore (\_).

## Enabling and disabling interoperability mode

Before enabling interoperability mode, inspect the individual fabrics for compatibility.

- Make sure that zones meet the zone criteria and restrictions listed in “[Zoning restrictions](#)” on page 221).
- Use the `msplmgmtdeactivate` command to disable any platform management functions. (Refer to the *HP StorageWorks Fabric OS 4.x command reference guide* for additional command information).

**To enable interoperability mode:**

1. Verify that you have implemented all the HP StorageWorks prerequisites necessary to enable interoperability mode on the fabric (see “[Configuration recommendations](#)” on page 221 and “[Configuration restrictions](#)” on page 221).
2. Connect to the switch and log in as admin.
3. Issue the `switchdisable` command to disable the switch.
4. Issue the `configure` command to set the domain ID to a number in the range from 97 to 126. (For detailed instructions, see “[Working with domain IDs](#)” on page 31.)
5. Issue the `interopmode 1` command to enable interoperability. This command resets a number of parameters and enables interactive mode.
6. Reboot the switch after changing the interoperability mode.

Example:

```
switch:admin> switchdisable
switch:admin> configure
Configure...
Fabric Parameters (yes, y, no, n): [no] y
Domain (1...239): [1] 97
...
switch:admin> interopmode 1
The switch effective configuration will be lost when the operating mode
is changed; do you want to continue? (yes, y, no, n): [no] y

done.
Interopmode is enabled

Note: It is recommended that you reboot this switch for the new change
to take effect.
switch:admin>
```

7. Repeat [step 2](#) through [step 6](#) on all HP StorageWorks switches in the fabric.
8. Enable interoperability on each non-HP StorageWorks switch. (Refer to the switch documentation.)
9. After enabling interoperability mode on all switches, physically connect the non-HP switches into the HP StorageWorks fabric, one at a time.

**To disable interoperability mode:**

1. Connect to the switch and log in as admin.
2. Issue the `switchdisable` command to disable the switch.
3. Issue the `interopmode 0` command to disable interoperability. This command resets a number of parameters and disables interactive mode.
4. Reboot the switch after changing the interoperability mode.

**Example:**

```
switch:admin> switchDisable
switch:admin> interopmode 0
The switch effective configuration will be lost when the operating mode
is changed; do you want to continue? (yes, y, no, n): [no] y

done.
Interopmode is disabled

Note: It is recommended that you reboot this switch for the new change to
take effect.
switch:admin>
```

- 5.** Wait for a fabric reconfiguration after removing each switch.
- 6.** Each non-HP StorageWorks switch may require the execution of a similar command to disable interoperability.
- 7.** Repeat this procedure on all HP StorageWorks switches in the fabric.

# C Using Remote Switch

The HP Remote Switch feature is available as part of the Fabric OS standard feature set through the use of the `portcfgislmode` command, which is described under “[Linking through a gateway](#)” on page 32. For those who use Remote Switch as part of their legacy set of tools, this appendix contains a description and procedure for the feature.

Remote Switch, enables you to connect two remote HP StorageWorks switches over an IP network, enabling communication of IP or ATM protocols as well as Fibre Channel traffic.

The HP Remote Switch feature functions with the aid of a bridging device or Fibre Channel gateway. The gateway supports both a Fibre Channel physical interface and a secondary, non-Fibre Channel physical interface, such as IP, SONET, or ATM. Remote Switch functions over E\_Port connections. With Remote Switch on both fabrics, the gateway accepts Fibre Channel frames from one fabric, tunnels them across the network, and passes them to the other fabric. From the viewpoint of the connected hosts and storage devices, fabrics using Remote Switch interact the same as locally connected switches.

Remote Switch provides many of the same capabilities of normal ISL links including,

- Coordinated fabric services

The Remote Switch fabric configuration fully supports all fabric services, including distributed name service, registered state change notification, and alias service.

- Distributed management

Management tools such as Advanced Web Tools, Fabric OS, and SNMP are available from both the local switch and the remote switch. Switch management is routed through the Fibre Channel connection; thus, no additional network connection is required between sites.

- Support for interswitch links (ISLs)

Sites requiring redundant configurations can connect multiple E\_Ports to remote sites by using multiple gateways. Standard Fabric OS routing facilities automatically maximize throughput and provide automatic failover during interruption on the WAN connection.

The Remote Switch feature operates in conjunction with a gateway. The gateway provides an E\_Port interface that links to the HP StorageWorks E\_Port. After the link between the two E\_Ports has been negotiated, the gateway E\_Port moves to passthrough mode and passes Fibre Channel traffic from the HP StorageWorks E\_Port to the WAN.

The gateway accepts Fibre Channel frames from one side of a Remote Switch fabric, transfers them across a WAN, and passes them to the other side of the Remote Switch fabric.

Remote Switch can be used for these types of gateway devices:

- Fibre Channel over ATM
- Fibre Channel over IP
- Fibre Channel over SONET
- Fibre Channel over DWDM

Most of these gateway devices have enough buffers to cover data transfer over a wide area network (WAN). The HP StorageWorks switches on each side of the gateway must have identical configurations. Only qualified SFPs should be used.

You must connect the fabrics through the gateway device, and make sure that the `configure` command parameters are compatible with the gateway device.

You may be required to reconfigure the following parameters, depending on the gateway requirements:

- `R_A_TOV`: Specify a Resource Allocation Timeout Value compatible with your gateway device.
- `E_D_TOV`: Specify a Error Detect Timeout Value compatible with your gateway device
- Data field size: Specify the maximum Fibre Channel data field reported by the fabric. Verify the maximum data field size the network-bridge can handle. Some bridges may not be able to handle a maximum data field size of 2112.
- BB credit: Specify the number of Buffer-to-Buffer credits for `Nx_port` devices.
- Suppress Class F Traffic: Use this parameter to disable class F traffic. Some network-bridge devices may not have a provision for handling class F frames. In this case, the transmission of class F frames must be suppressed throughout the entire Remote Switch fabric.

**To set the access and reconfigure these parameters:**

1. Connect to the switch and log in as admin.
2. Issue the `switchdisable` command to disable the switch.
3. Issue the `configure` command.
4. Enter `yes` at the `Fabric Parameters` prompt.
5. Press **Enter** to scroll through the Fabric Parameters without changing their values, until you reach the parameter you want to modify.
6. Specify a new parameter value that is compatible with your gateway device.
7. Press **Enter** to scroll through the remainder of the configuration parameters. Make sure that the configuration changes are committed to the switch.
8. Repeat for all switches in the fabrics to be connected through a gateway device. These parameters must be identical on each switch in the fabric, and between fabrics connected through the gateway device.

The following example shows how to modify the data field size and suppress class F traffic on a switch:

```
switch:admin> switchDisable
switch:admin> configure
Configure...
Fabric parameters (yes, y, no, n): [no] yes
Domain: (1..239) [3]
    R_A_TOV: (4000..120000) [10000]
    E_D_TOV: (1000..5000) [2000]
    Data field size: (256..2112) [2112] 1000
    Sequence Level Switching: (0..1) [0]
    Disable Device Probing: (0..1) [0]
    Suppress Class F Traffic: (0..1) [0] 1
    VC Encoded Address Mode: (0..1) [0]
    Per-frame Route Priority: (0..1) [0]
    Long Distance Fabric: (0..1) [0]
    BB credit: (1..16) [16]
Virtual Channel parameters (yes, y, no, n): [no]
    Zoning Operation parameters (yes, y, no, n): [no]
    RSCN Transmission Mode (yes, y, no, n): [no]
    NS Operation Parameters (yes, y, no, n): [no]
    Arbitrated Loop parameters (yes, y, no, n): [no]
    System services (yes, y, no, n): [no]
    Portlog events enable (yes, y, no, n): [no]
Committing configuration...done.
switch:admin>
```



# D Understanding legacy password behavior

This appendix contains the following sections:

- [Password management information](#), page 229
- [Password prompting behaviors](#), page 232
- [Password migration during firmware changes](#), page 233
- [Password recovery options](#), page 235

The following sections provide password information for early versions of Fabric OS firmware.

## Password management information

[Table 39](#) describes the password standards and behaviors between various versions of firmware.

**Table 39** Account and password characteristics matrix

Topic	V2.6.0/3.0.0	V2.6.2/3.1.0	V4.0.0	V4.1.0 to v4.2.0
Number of accounts on the switch	4	4	4, chassis based	Core Switch 2/64: 8 for the chassis, 4 per switch  SAN Switch 2/8V, 2/16V, 2/32, 4/32, SAN Director 2/128
Account login names	root, factory, admin, user	root, factory, admin, user	root, factory, admin, user	root, factory, admin, user. 3016: administrative account is called USERID.
Account name changing feature	Yes, when Secure FabOS is disabled; No, when Secure FabOS is enabled.	Yes, when Secure FabOS is disabled; No, when Secure FabOS is enabled.	No	No, regardless of security mode.
Maximum and minimum amount of characters for a password	8 - 40 characters with printable ASCII	8 - 40 characters with printable ASCII	0 - 8 (Standard UNIX)	8 - 40 characters with printable ASCII

**Table 39** Account and password characteristics matrix (continued)

Topic	V2.6.0/3.0.0	V2.6.2/3.1.0	V4.0.0	V4.1.0 to v4.2.0
Can different switch instances use a different password for the same account login level? For example, the password for admin for switch 0 can be different from password for admin for switch 1.	n.a.	n.a.	No	Yes for the Core Switch 2/64. N/A for all other switches.
Does the root account use restricted shell?	No	No	No	No
When connecting to a factory installed switch, do you use the default passwords?	Yes	Yes	Yes	Yes
Does a user need to know the old passwords when changing passwords using the <code>passwd</code> command?	Yes, the old password is required to change any password, regardless of the level at which you connect.	Yes, the old password is required to change any password, regardless of the level at which you connect.	Yes, except when the root user changes another user's password. This is standard UNIX behavior; Fabric OS does not enforce any additional security.	Old password is required only when changing password for the same level user password. Changing password for lower level user does not require old password. For example, users connect as admin; old admin password is required to change the admin password. But old user password is not required to change the user password.
Can <code>passwd</code> change higher-level passwords? For example, can admin change root password?	No. If users connect as admin, the users can only change admin and user passwords; the users cannot change factory, nor root password.	No. If users connect as admin, the users can only change admin and user passwords; the users cannot change factory, nor root password.	Yes, but asks for the old password of the higher-level account (example "root").	Yes; if users connect as admin, they can change the root, factory, and admin passwords. However, if one connects as user, one can only change the user password.

**Table 39** Account and password characteristics matrix (continued)

Topic	V2.6.0/3.0.0	V2.6.2/3.1.0	V4.0.0	V4.1.0 to v4.2.0
Can API change passwords?	API can change admin passwords on any switch, when security mode is disabled. It can only change the admin password on the primary FCS switch when security mode is enabled.	API can change admin passwords on any switch, when security mode is disabled. It can only change the admin password on the primary FCS switch when security mode is enabled.	Yes, only for admin.	Yes, only for admin.
Can Advanced Web Tools change passwords?	When security mode is disabled, users can change the admin and user passwords on all switches using Advanced Web Tools. When security mode enabled, users can only change the admin and user passwords on the primary FCS switch using Advanced Web Tools.	When security mode is disabled, users can change the admin and user passwords on all switches using Advanced Web Tools. When security mode enabled, users can only change the admin and user passwords on the primary FCS switch using Advanced Web Tools.	No	No
Can SNMP change passwords?	No	No	No	No

# Password prompting behaviors

Table 40 describes the expected password prompting behaviors of various Fabric OS versions.

**Table 40** Password prompting matrix

Topic	V2.6.0/3.0.0	V2.6.2/3.1.0	V4.0.0	V4.1.0 and later
Must all password prompts be completed for any change to take effect?	Yes. If users only provide some of the passwords before exiting, no passwords are changed. Prompting continues on the next appropriate login.	Yes. If users only provide some of the passwords before exiting, no passwords are changed. Prompting continues on the next appropriate login.	No. Partial changes of all four passwords are allowed.	No. Partial changes of all four passwords are allowed.
When does the password prompt appear?	When users connect as root, factory, or admin.	When users connect as root, factory, or admin.	When users connect as root, factory, or admin, the accounts with default password are prompted for change. The accounts with non-default password are not be prompted.	When users connect as root, factory, or admin, the accounts with default password are prompted for change. The accounts with non-default password are not be prompted.
Is a user forced to answer password prompts before getting access to the firmware?	No, users can type in <b>Ctrl+c</b> to get out of password prompting.	No, users can type in <b>Ctrl+c</b> to get out of password prompting.	No, users can type in <b>Ctrl+c</b> to get out of password prompting.	No, users can type in <b>Ctrl+c</b> to get out of password prompting.
Do users need to know the old root password when answering prompting?	No	No	Yes in v4.0.0 *No in v4.0.2 only	No
Are new passwords forced to be set to something different than the old passwords?	Yes	Yes	Yes	Yes
Is password prompting disabled when security mode is enabled?	Yes	Yes	Yes	Yes

**Table 40** Password prompting matrix (continued)

<b>Topic</b>	<b>V2.6.0/3.0.0</b>	<b>V2.6.2/3.1.0</b>	<b>V4.0.0</b>	<b>V4.1.0 and later</b>
Is the passwd command disabled until the user has answered password prompting?	True	True	False	True
Does password prompting reappear when passwords are changed back to default using the passwd command?	No	No	Yes	No
Does password prompting reappear when passwords are changed back to default using the passwdDefault command?	Yes	Yes	Yes	Yes.

## Password migration during firmware changes

Table 41 describes the expected outcome of password settings when upgrading or downgrading firmware for various Fabric OS versions.

**Table 41** Password migration behavior during firmware upgrade and downgrade

<b>Topic</b>	<b>V2.4.0/2.6.0</b>	<b>V3.0.0/3.1.0</b>	<b>V4.0.0</b>	<b>V4.1.0 to v4.4.0</b>
Passwords used when upgrading to a newer firmware release for the first time.	For first time firmware upgrades from v2.4.0 to v2.6.0, the v2.4.0 passwords are preserved.	For first time firmware upgrades from v3.0.0 to v3.1.2, the v3.0.0 passwords are preserved.	n.a.	For first time firmware upgrades from v4.0.0 to v4.2.0, the v4.0.0 passwords are preserved.
Passwords preserved during subsequent firmware upgrades	For second firmware upgrades (and each subsequent upgrade) from v2.4.0 to v2.6.0, the passwords that were last used in v2.6.0 are effective.	For second firmware upgrades (and each subsequent upgrade) from v3.0.0 to v3.1.0, the passwords that were last used in v3.1.0 are effective.	n.a.	For second firmware upgrades (and each subsequent upgrade) from v4.0.0 to v4.2.0 the passwords that were last used in v4.0.0 are effective.

**Table 41** Password migration behavior during firmware upgrade and downgrade (continued)

Topic	V2.4.0/2.6.0	V3.0.0/3.1.0	V4.0.0	V4.1.0 to v4.4.0
Is downgrading to an older firmware version (which does not support Secure Fabric OS) allowed when security mode is enabled?	Yes. FirmwareDownload does not prevent such downgrades.	Yes	n.a.	Yes
Passwords used if downgrading to an older firmware for the first time	When downgrading firmware from v2.6.0 to v2.4.0 for the first time, the default passwords are used.	When downgrading firmware from v3.1.2 to v3.0.0 for the first time, the default passwords are used.	n.a.	If the switch had v4.2.0 factory installed, a firmware downgrade from v4.2.0 to v4.0.0 uses the default passwords.
When downgrading to an older firmware at subsequent times, which passwords are used?	Firmware downgrades from v2.6.0 to v2.4.0 use the previous v2.4.0 passwords (the passwords used before the firmware had been upgraded to v2.6.0).	Firmware downgrades from v3.1.0 to v3.0.0 use the previous v3.0.0 passwords (the passwords used before the firmware had been upgraded to v3.1.0).	Firmware downgrades within 4.0.0 use the old 4.0.0 passwords.	Firmware downgrades from v4.2.0 to v4.0.0 use the previous v4.0.0 passwords (the passwords used before the firmware had been upgraded to v4.2.0).
When downgrading then upgrading again, what passwords are used?	When upgrading firmware for a second time, the old v2.6.0 or v3.1.0 passwords are used (the passwords used before the firmware had been downgraded).	When upgrading firmware for a second time, the old v2.6.0 or v3.1.0 passwords are used (the passwords used before the firmware had been downgraded).	When upgrading firmware for a second time, the old passwords are used (the passwords used before the firmware had been downgraded).	When upgrading firmware for a second time, the old passwords are used (the passwords used before the firmware had been downgraded). For a 4.0.0 to 4.2.0, use the 4.0.0 passwords

# Password recovery options

Table 42 describes the options available when one or more types of passwords are lost.

**Table 42** Password recovery options

Topic	V2.6.0/3.0.0	V3.0.0.2/3.1.0	V4.0.0	V4.1.0 and later
If all the passwords are forgotten, what is the password recovery mechanism? Are these procedures non-disruptive recovery procedures?	The user has to get a special password recovery firmware based on the WWN of the switch from Tech Support, and then download the special firmware; this resets all passwords to default. The procedures are disruptive.	The user has to get a special password recovery firmware based on the WWN of the switch from Tech Support, and then download the special firmware; this resets all passwords to default. The procedures are disruptive.	Contact your switch service provider. A non-disruptive procedure is available.	Contact your switch service provider. A non-disruptive procedure is available.
If a user has only the root password, what is the password recovery mechanism?	Use passwdDefault command to set all passwords to default.	Use passwdDefault command to set all passwords to default.	Root can change any password by using the passwd command.	Use passwd command to set other passwords.  Use passwdDefault command to set all passwords to default.
How to recover boot PROM password?	n.a.	n.a.	n.a.	Contact your switch service provider and provide the recovery string.  See " <a href="#">Setting the boot PROM password</a> " on page 70 for instructions on setting the password with a recovery string.
How do I recover a user, admin, or factory password?	See " <a href="#">Recovering forgotten passwords</a> " on page 73.			



# E Zone merging scenarios

Table 43 provides information on merging zones and the expected results.

**Table 43** Zone merging scenarios

Description	Switch A	Switch B	Expected results
Switch A with a defined configuration Switch B does not have a defined configuration	defined: cfg1: zone1: ali1; ali2 effective: none	defined: none effective: none	Configuration from Switch A to propagate throughout the fabric in an inactive state, because the configuration is not enabled.
Switch A with a defined and enabled configuration Switch B has a defined configuration but no effective configuration	defined: cfg1 zone1: ali1; ali2 effective: cfg1:	defined: cfg1 zone1: ali1; ali2 effective: none	Configuration from Switch A to propagate throughout the fabric.
Switch A and Switch B have the same defined configuration. Neither have an enabled configuration.	defined: cfg1 zone1: ali1; ali2 effective: none	defined: cfg1 zone1: ali1; ali2 effective: none	No change (clean merge).
Switch A and Switch B have the same defined and enabled configuration.	defined: cfg1 zone1: ali1; ali2 effective: cfg1:	defined: cfg1 zone1: ali1; ali2 effective: cfg1:	No change (clean merge).
Switch A does not have a defined configuration Switch B with a defined configuration	defined: none effective: none	defined: cfg1 zone1: ali1; ali2 effective: none	Switch A absorbs the configuration from the fabric.
Switch A does not have a defined configuration Switch B with a defined configuration	defined: none effective: none	defined: cfg1 zone1: ali1; ali2 effective: cfg1	Switch A absorbs the configuration from the fabric. With cfg1 as the effective cfg.
Switch A and Switch B have the same defined configuration. Only Switch B has an enabled configuration.	defined: cfg1 zone1: ali1; ali2 effective: none	defined: cfg1 zone1: ali1; ali2 effective: cfg1	Clean merge - with cfg1 as the effective cfg.
Switch A and Switch B have different defined configurations. Neither have an enabled zone configuration.	defined: cfg2 zone2: ali3; ali4 effective: none	defined: cfg1 zone1: ali1; ali2 effective: none	Clean merge - the new cfg is a composite of the two - defined: cfg1 zone1: ali1; ali2 cfg2: zone2: ali3; ali4 effective: none

**Table 43** Zone merging scenarios (continued)

Description	Switch A	Switch B	Expected results
Switch A and Switch B have different defined configurations. Switch B has an enabled configuration.	defined: cfg2 zone2: ali3; ali4 effective: none	defined: cfg1 zone1: ali1; ali2 effective: cfg1	Clean merge - The new cfg is a composite of the two, with cfg1 as the effective cfg.
Effective cfg mismatch	defined: cfg1 zone1: ali1; ali2 effective: cfg1 zone1: ali1; ali2	defined: cfg2 zone2: ali3; ali4 effective: cfg2 zone2: ali3; ali4	Fabric segments due to: Zone Conflict cfg mismatch
cfg content mismatch	defined: cfg1 zone1: ali1; ali2 effective: irrelevant	defined: cfg1 zone1: ali3; ali4 effective: irrelevant	Fabric segments due to: Zone Conflict content mismatch
	defined: cfg1 zone1: ali1; ali2 effective: irrelevant	defined: cfg1 zone1: ali1; ali4 effective: irrelevant	Fabric segments due to: Zone Conflict content mismatch
Same content - different effective cfg name	defined: cfg1 zone1: ali1; ali2 effective: cfg1 zone1: ali1; ali2	defined: cfg2 zone1: ali1; ali2 effective: cfg2 zone1: ali1; ali2	Fabric segments due to: Zone Conflict cfg mismatch
Same content - different zone name	defined: cfg1 zone1: ali1; ali2 effective: irrelevant	defined: cfg1 zone2: ali1; ali2 effective: irrelevant	Fabric segments due to: Zone Conflict content mismatch
Same content - different alias name	defined: cfg1 ali1: A; B effective: irrelevant	defined: cfg1:ali2: A; B effective: irrelevant	Fabric segments due to: Zone Conflict content mismatch
Same name - different types	effective: zone1: MARKETING	effective: cfg1: MARKETING	Fabric segments due to: Zone Conflict type mismatch
Same name - different types	effective: zone1: MARKETING	effective: alias1: MARKETING	Fabric segments due to: Zone Conflict type mismatch
Same name - different types	effective:cfg1: MARKETING	effective:alias1: MARKETING	Fabric segments due to: Zone Conflict type mismatch

# F Upgrading firmware in single mode

In the SAN Switch 2/32, the Core Switch 2/64, and the SAN Director 2/128, the `firmwareDownload` command, by default, performs a full installation, automatic reboot (`autoreboot`), and automatic firmware commit (`autocommit`). Automatic reboot and automatic commit modes are not selectable by default; however, they become selectable when *single mode* is enabled by entering the `-s` option on the command line. In this case, `firmwareDownload` disables autoreboot and enables autocommit.



**NOTE:** You should only use the following procedures if instructed to do so by your service provider.

Your service provider may ask you to perform this procedure on the SAN Switch 2/32 under these circumstances:

- To prevent the firmware commit that occurs after downloading, so that you can restore previous versions.
- To control the timing of the execution of the `hareboot` command, so that you can prestage the firmware ahead of time.

Your service provider may ask you to perform the single CP card procedure on the Core Switch 2/64 and the SAN Director 2/128 if a CP card fails and the replacement CP card is running a version of firmware that cannot synchronize with the current active CP card.

For information about messages that may appear during the procedures, refer to the *HP StorageWorks Fabric OS 4.x diagnostics and system error messages reference guide*.

To upgrade HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32 in single mode:

Specify the `-s` option on the command line to enable the selection of autocommit and autoreboot:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
firmwareDownload -s
```

Make sure to type a space between the command and the `-s` option.

3. Enter the IP address of the FTP server where the firmware is stored.
4. Enter your user name for the server.
5. Enter the full path to the firmware file on the server; for example:

```
/pub/v4.4.0/release.plist
```

6. Enter your password.

**7.** Answer the next prompts as follows:

Do Auto Commit after reboot [Y]: **y**

If you specify no, you must manually enter the `firmwarecommit` command.

Reboot system after download [N]: **y**

The default is no. If you take the default, you must later use the `hareboot` command to perform a high-availability reboot manually.

Full Install (Otherwise upgrade only) [Y]: **y**

After you upgrade to v4.4.0 or later, this option is no longer available.

**8.** Wait for the firmware download to finish.

**9.** Start a new telnet session and use the `firmwaredownloadstatus` command to check the status.

**Example**

```
switch: admin> firmwaredownload -s
Server Name or IP Address: 192.1.2.3
User Name: JohnDoe
File Name: /pub/v4.4.0/release.plist
Password: *****
Full Install (Otherwise upgrade only) [Y]: y
Do Auto Commit after reboot [Y]: y
Reboot system after download [N]: y
Firmwaredownload has started.

.
```

**To upgrade a single Core Switch 2/64 or SAN Director 2/128 CP card:**

Although it is possible to upgrade firmware on one CP card at a time, you should not do so under normal circumstances because it disrupts a functioning fabric.

When the two CP cards are not running the same firmware versions, it may be necessary to disable one or the other to maintain fabric stability. For information on the commands used to achieve this, refer to the `hadisable` and `hafailover` commands in the *HP StorageWorks Fabric OS 4.x command reference guide*.

The following procedure allows you to upgrade a single CP card. This procedure can be used with Fabric OS v4.0.0d and later.

**To upgrade to a single CP card:**

1. Connect to the switch and log in as admin.
2. Issue the `hashow` command to determine which CP card is the active and which one is the standby. In the following example, the active CP card is CP0, and the standby CP card is CP1:

```
switch:admin> hashow
Local CP (Slot 5, CP0): Active
Remote CP (Slot 6, CP1): Standby, Healthy
HA enabled, Heartbeat Up, HA State in sync
```

3. Log in to the standby CP card as admin.

4. Issue the following command to upgrade a new version of the firmware to the standby CP card:

`firmwaredownload -s`

Make sure to type a space between the command and the `-s` option.

5. Enter the IP address of the FTP server where the firmware is stored.
6. Enter your user name for the server.
7. Enter the full path to the firmware file on the server; for example:  
`/pub/v4.4.0/release.plist`
8. Enter your password.
9. Answer the next prompts as indicated:

`Do Auto Commit after reboot [Y]: y`

If you answer no, you must manually enter the `firmwarecommit` command.

`Reboot system after download [N]: y`

The default is no. If you take the default, you must later use the `hareboot` command to perform a high-availability reboot manually.

`Full Install (Otherwise upgrade only) [Y]: y`

After you upgrade to v4.4.0 or later, this option is no longer available.

The default is no. If you take the default, you must manually reboot the system.

Example:

```
switch: admin> firmwaredownload -s
Server Name or IP Address: 192.1.2.3
User Name: JohnDoe
File Name: /pub/v4.4.0/release.plist
Password: *****
Full Install (Otherwise upgrade only) [Y]: y
Do Auto Commit after reboot [Y]: y
Reboot system after download [N]: y
```

10. Wait for the firmware download to finish.
11. Start a new telnet session and use the `firmwaredownloadstatus` command to check the status.)
12. Issue the `hashow` command to verify that the two CP cards are synchronized.
13. Reboot the standby CP card (if you set the option to reboot automatically to no in [step 9](#)).
14. Log in to the same CP card and enter the `firmwaredownloadstatus` command to verify firmware has downloaded successfully and has either committed or is in the process of doing so. (If you set the option to do auto commit after reboot to no in [step 9](#), you must enter the `firmwarecommit` command manually.)
15. Issue the `hafailover` command to make the current CP card active (with the updated firmware).
16. Repeat [step 1](#) through [step 14](#) on the other CP card if the firmware versions are different.



# Glossary

## **8b/10b encoding**

An encoding scheme that converts each 8-bit byte into 10 bits. Used to balance 1s and 0s in high-speed transports.

## **ABTS**

Abort Basic Link Service. Also called an *Abort Sequence*.

## **ACC**

Accept link service reply. The normal reply to an Extended Link Service request (such as FLOGI), indicating that the request has been completed.

## **address identifier**

A 24-bit or 8-bit value used to identify the source or destination of a frame. See also [S\\_ID](#) and [D\\_ID](#).

## **AL\_PA**

Arbitrated loop physical address. A unique 8-bit value assigned during loop initialization to a port in an arbitrated loop. Alternately, *arbitrated loop parameters*.

## **AL\_TIME**

Arbitrated loop time-out value. Twice the amount of time it would take for a transmission word to propagate around a worst-case loop. The default value is 15 milliseconds.

## **alias**

A logical grouping of elements in a fabric. An alias is a collection of port numbers and connected devices used to simplify the entry of port numbers and WWNs when creating zones.

## **alias address identifier**

An address identifier recognized by a port in addition to its standard identifier. An alias address identifier can be shared by multiple ports. See also [alias](#).

## **alias AL\_PA**

An AL\_PA value recognized by an L\_Port in addition to the AL\_PA assigned to the port. See also [AL\\_PA](#).

## **alias server**

A fabric software facility that supports multicast group management.

## **ARB**

Arbitrative primitive signal. Applies only to an arbitrated loop topology. Transmitted as the fill word by an L\_Port to indicate that the port is arbitrating access to the loop.

## **arbitrated loop**

A shared 100-MBps Fibre Channel transport structured as a loop. Can support up to 126 devices and one fabric attachment. See also [topology](#).

**arbitration**

A method of gaining orderly access to a shared-loop topology.

**area number**

In Fabric OS v4.0.0 and later, ports on a switch are assigned a logical area number. Port area numbers can be viewed by use of the `switchshow` command. Area numbers define the operative port for many Fabric OS commands; for example, area numbers can be used to define the ports within an alias or zone.

**ARR**

Asynchronous response router. Refers to Management Server GS\_Subtype Code E4, which appears in `portlogdump` command output.

**ASD**

Alias server daemon. Used for managing multicast groups by supporting the create, add, remove, and destroy functions.

**ATM**

Asynchronous Transfer Mode. A transport used for transmitting data over LANs or WANs that transmit fixed-length units of data. Provides any-to-any connectivity and allows nodes to transmit simultaneously.

**authentication**

The process of verifying that an entity in a fabric (such as a switch) is what it claims to be. See also [digital certificate](#), [switch-to-switch authentication](#).

**autocommit**

A feature of the `firmwaredownload` command. Enabled by default, `autocommit` commits new firmware to both partitions of a control processor.

**autoreboot**

Refers to the `-b` option of the `firmwaredownload` command. Enabled by default.

**BB\_Credit**

Buffer-to-buffer credit. The number of frames that can be transmitted to a directly connected recipient or within an arbitrated loop. Determined by the number of receive buffers available. See also [buffer-to-buffer flow control](#).

**beacon**

A tool in which all of the port LEDs on a switch are set to flash from one side of the switch to the other, to enable identification of an individual switch in a large fabric. A switch can be set to beacon by a CLI command or through Advanced Web Tools.

**beginning running disparity**

The disparity at the transmitter or receiver when the special character associated with an ordered set is encoded or decoded. See also [disparity](#).

**BIST**

Built-in self-test.

**bit synchronization**

The condition in which a receiver is delivering retimed serial data at the required bit error rate.

**block**

As applied to Fibre Channel technology, upper-level application data that is transferred in a single sequence.

**broadcast**

The transmission of data from a single source to all devices in the fabric, regardless of zoning. See also [multicast](#), [unicast](#).

**buffer-to-buffer flow control**

Management of the frame transmission rate in either a point-to-point topology or in an arbitrated loop. See also [BB\\_Credit](#).

**bypass circuitry**

Circuits that automatically remove a device from the data path when valid signals are dropped.

**CA**

Certificate authority. A trusted organization that issues digital certificates. See also [digital certificate](#).

**CAM**

Content-addressable memory.

**Class 1 service**

The class of frame-switching service for a dedicated connection between two communicating ports (also called *connection-oriented service*). Includes acknowledgement of frame delivery or nondelivery.

**Class 2 service**

A connectionless class of frame-switching service that includes acknowledgement of frame delivery or nondelivery.

**Class 3 service**

A connectionless class of frame-switching service that does not include acknowledgement of frame delivery or nondelivery. Can be used to provide a multicast connection between the frame originator and recipients, with acknowledgement of frame delivery or nondelivery.

**Class 4 service**

A connection-oriented service that allows fractional parts of the bandwidth to be used in a virtual circuit.

**Class 6 service**

A connection-oriented multicast service geared toward video broadcasts between a central server and clients.

**Class F service**

The class of frame-switching service for a direct connection between two switches, allowing communication of control traffic between the E\_Ports. Includes acknowledgement of data delivery or nondelivery.

**class of service**

A specified set of delivery characteristics and attributes for frame delivery.

**CLS**

Close primitive signal. Used only in an arbitrated loop. Sent by an L\_Port that is currently communicating in the loop, to close communication to another L\_Port.

## **configuration**

1. A set of parameters that can be modified to fine-tune the operation of a switch. Use the `configshow` command to view the current configuration of your switch.
2. In HP Zoning, a zoning element that contains a set of zones. The configuration is the highest-level zoning element and is used to enable or disable a set of zones on the fabric. See also [zone configuration](#).

## **COS**

Class of service.

## **CP**

Control processor.

## **credit**

As applied to Fibre Channel technology, the number of receive buffers available to transmit frames between ports. See also [BB\\_Credit](#).

## **D\_ID**

Destination identifier. A three-byte field in the frame header that indicates the address identifier of the N\_Port to which the frame is headed.

## **defined zone configuration**

The set of all zone objects defined in the fabric. Can include multiple zone configurations. See also [zone configuration](#).

## **digital certificate**

An electronic document issued by a certificate authority (CA) to an entity, containing the public key and identity of the entity. Entities in a secure fabric are authenticated based on these certificates. See also [authentication](#), [CA](#), [public key](#).

## **disparity**

The proportion of 1s and 0s in an encoded character. *Neutral disparity* means an equal number of each; *positive disparity* means a majority of 1s; *negative disparity* means a majority of 0s.

## **DLS**

Dynamic load-sharing. Dynamic distribution of traffic over available paths. Allows for recomputing of routes when an Fx\_Port or E\_Port changes status.

## **domain controller**

A domain controller (or embedded port) communicates with and gets updates from other switches' embedded ports. The well-known address is `ffffcdd`, where `dd` is the domain number).

## **domain ID**

A unique identifier for all switches in a fabric, used in routing frames. Usually assigned by the principal switch but can be assigned manually. The domain ID for an HP StorageWorks switch can be any integer from 1 through 239.

## **E\_D\_TOV**

Error-detect time-out value. The minimum time a target waits for a sequence to complete before initiating recovery. Can also be defined as the maximum time allowed for a round-trip transmission before an error is declared. See also [R\\_A\\_TOV](#), [RR\\_TOV](#).

## **E\_Port**

Expansion port. A type of switch port that can be connected to an E\_Port on another switch to create an ISL. See also [ISL](#).

## **ELP**

Exchange link parameters.

## **ELS**

Extended link service. ELSs are sent to the destination N\_Port to perform the requested function or service. ELS is a Fibre Channel standard that is sometimes called a *Fibre Channel Physical (FC\_PH) ELS*.

## **EM**

Environmental monitor. Monitors FRUs and reports failures.

## **embedded port**

An embedded port (or domain controller) communicates and gets updates from other switches' embedded ports. The well-known address is *ffffcdd*, where *dd* is the domain number.

## **entry fabric**

The basic HP software license that allows one E\_Port per switch.

## **EOF**

End of frame. A group of ordered sets used to mark the end of a frame.

## **error**

As applied to the Fibre Channel industry, a missing or corrupted frame, time-out, loss of synchronization, or loss of signal (link errors).

## **exchange**

The highest-level Fibre Channel mechanism used for communication between N\_Ports. Composed of one or more related sequences, it can work in either one or both directions.

## **F\_BSY**

Fabric port busy frame. A frame issued by the fabric to indicate that a frame cannot be delivered because the fabric or destination N\_Port is busy.

## **F\_Port**

Fabric port. A port that is able to transmit under fabric protocol and interface over links. Can be used to connect an N\_Port to a switch. See also [FL\\_Port](#), [Fx\\_Port](#).

## **F\_RJT**

Fabric port reject frame. A frame issued by the fabric to indicate that delivery of a frame is being denied, either because a class is not supported, the existence of an invalid header, or no N\_Port available.

## **fabric**

A Fibre Channel network containing two or more switches in addition to hosts and devices. Also called a *switched fabric*. See also [SAN](#), [topology](#).

## **Fabric Manager**

Optionally licensed HP software. Fabric Manager is a GUI that allows for fabric-wide administration and management. Switches can be treated as groups, and actions such as firmware downloads can be performed simultaneously.

### **fabric name**

The unique identifier assigned to a fabric and communicated during login and port discovery.

### **fabric services**

Codes that describe the communication to and from any well-known address.

### **fabric topology**

The arrangement of switches that form a fabric.

### **Fabric Watch**

Optionally licensed HP software. Fabric Watch can be accessed through either the command line or Advanced Web Tools; it provides the ability to set thresholds for monitoring fabric conditions.

### **failover**

Describes the Core Switch 2/64 process of one CP card passing active status to another CP card. A failover is nondisruptive.

### **FAN**

Fabric address notification. Retains the AL\_PA and fabric address when a loop reinitializes, if the switch supports FAN.

### **FC-0**

Lowest layer of Fibre Channel transport. Represents physical media.

### **FC-1**

Layer of Fibre Channel transport that contains the 8b/10b encoding scheme.

### **FC-2**

Layer of Fibre Channel transport that handles framing and protocol, frame format, sequence and exchange management, and ordered set usage.

### **FC-3**

Layer of Fibre Channel transport that contains common services used by multiple N\_Ports in a node.

### **FC-4**

Layer of Fibre Channel transport that handles standards and profiles for mapping upper-level protocols such as SCSI and IP onto the Fibre Channel Protocol.

### **FC-CT**

Fibre Channel common transport.

### **FC-FG**

Fibre Channel generic requirements.

## **FC-GS**

Fibre Channel generic services.

## **FC-GS-2**

Fibre Channel generic services, second generation.

## **FC-GS-3**

Fibre Channel Generic Services, third generation.

## **FC\_IP**

Fibre Channel-Over-IP.

## **FC-PH**

The Fibre Channel physical and signaling standard for FC-0, FC-1, and FC-2 layers of the Fibre Channel Protocol. Indicates signaling used for cable plants, media types, and transmission speeds.

## **FCP**

Fibre Channel Protocol. Mapping of protocols onto the Fibre Channel standard protocols. For example, SCSI FCP maps SCSI-3 onto Fibre Channel.

## **FCS**

Fibre Channel Standard.

## **FCS switch**

With respect to the HP Secure Fabric OS feature, one or more designated switches that store and manage security parameters and configuration data for all switches in the fabric. They also act as a set of backup switches to the primary FCS switch. See also [primary FCS switch](#).

## **FC-SW-2**

The second-generation Fibre Channel Switch Fabric standard defined by ANSI. Specifies tools and algorithms for the interconnection and initialization of Fibre Channel switches to create a multiswitch Fibre Channel fabric.

## **FDMI**

Fabric-Device Management Interface. A database service provided by the fabric for Nx\_Ports. The primary use is by HBA devices that register information about themselves and their ports.

## **FFFFF5**

Well-known Fibre Channel address for a Class 6 multicast server.

## **FFFFF6**

Well-known Fibre Channel address for a clock synchronization server.

## **FFFFF7**

Well-known Fibre Channel address for a security key distribution server.

## **FFFFF8**

Well-known Fibre Channel address for an alias server.

**FFFFF9**

Well-known Fibre Channel address for a QoS facilitator.

**FFFFFA**

Well-known Fibre Channel address for a management server.

**FFFFFB**

Well-known Fibre Channel address for a time server.

**FFFFFC**

Well-known Fibre Channel address for a directory server.

**FFFFFD**

Well-known Fibre Channel address for a fabric controller.

**FFFFFE**

Well-known Fibre Channel address for a fabric F\_Port.

**FFFFFF**

Well-known Fibre Channel address for a broadcast alias ID.

**Fibre Channel**

A protocol that transmits data between servers, switches, and storage devices. It is a high-speed, serial, bidirectional, topology-independent, multiprotocol, and highly scalable interconnection between computers, peripherals, and networks.

**FIFO**

First in, first out. Refers to a data buffer that follows the first in, first out rule.

**firmware**

The basic operating system provided with the hardware.

**FL\_Port**

Fabric loop port. A port that is able to transmit under fabric protocol and also has arbitrated loop capabilities. Can be used to connect an NL\_Port to a switch. See also [F\\_Port](#), [Fx\\_Port](#).

**flash**

Programmable nonvolatile RAM (NVRAM) memory that maintains its contents without power.

**FLOGI**

Fabric login. The process by which an N\_Port determines whether a fabric is present and, if so, exchanges service parameters with it. See also [PLOGI](#).

**frame**

The Fibre Channel structure that transmits data between ports. Consists of a start-of-frame delimiter, header, optional headers, data payload, cyclic redundancy check (CRC), and end-of-frame delimiter. There are two types of frames: *link control frames* (transmission acknowledgements and so forth) and *data frames*.

**FRU**

Field-replaceable unit. A component that can be replaced onsite.

**FSPF**

Fabric shortest path first. The standard routing protocol for Fibre Channel switches.

**FSS**

Fabric OS state synchronization. The FSS service is related to HA. The primary function of FSS is to deliver state update messages from active components to their peer standby components. FSS determines whether fabric elements are synchronized (and thus FSS compliant).

**FTP**

File Transfer Protocol.

**full fabric**

The HP software license that allows multiple E\_Ports on a switch, making it possible to create multiple ISL links.

**full fabric citizenship**

A loop device that has an entry in the Simple Name Server.

**full-duplex**

A mode of communication that allows the same port to simultaneously transmit and receive frames. See also [half-duplex](#).

**Fx\_Port**

A fabric port that can operate as either an F\_Port or FL\_Port. See also [F\\_Port](#), [FL\\_Port](#).

**G\_Port**

Generic port. A port that can operate as either an E\_Port or an F\_Port. A port is defined as a G\_Port when it is not yet connected or has not yet assumed a specific function in the fabric.

**gateway**

Hardware that connects incompatible networks by providing translation for both hardware and software. For example, an ATM gateway can be used to connect a Fibre Channel link to an ATM connection.

**GBIC**

Gigabit interface converter. A removable serial transceiver module that allows gigabaud physical-level transport for Fibre Channel and gigabit Ethernet.

**Gbps**

Gigabits per second (1,062,500,000 bits/second).

**GBps**

Gigabytes per second (1,062,500,000 bytes/second).

## **GLM**

Gigabit Link Module. A semitransparent transceiver that incorporates serializing and deserializing functions.

## **GMT**

Greenwich Mean Time. An international time zone. Also known as *UTC*. See also [UTC](#).

## **GUI**

A graphical user interface, such as HP Advanced Web Tools and HP Fabric Manager.

## **HA**

High availability. A set of features in HP StorageWorks switches that provides maximum reliability and nondisruptive replacement of key hardware and software modules.

## **half-duplex**

A mode of communication that allows a port to either transmit or receive frames at any time except simultaneously (with the exception of link control frames, which can be transmitted at any time). See also [full-duplex](#).

## **hard address**

The AL\_PA that an NL\_Port attempts to acquire during loop initialization.

## **HBA**

Host bus adapter. The interface card between a server or workstation bus and the Fibre Channel network.

## **header**

A Fibre Channel frame has a header and a payload. The header contains control and addressing information associated with the frame. See also [payload](#).

## **HiPPI**

High-Performance Parallel Interface. An 800 Mbps interface normally used in supercomputer environments.

## **hop count**

The number of ISLs a frame must traverse to get from its source to its destination.

## **host**

A computer system that provides end users with services like computation and storage access.

## **HTTP**

Hypertext Transfer Protocol. The standard TCP/IP transfer protocol used on the World Wide Web.

## **hub**

A Fibre Channel wiring concentrator that collapses a loop topology into a physical star topology. Nodes are automatically added to the loop when active and removed when inactive.

## **idle**

Continuous transmission of an ordered set over a Fibre Channel link (when no data is being transmitted) to keep the link active and maintain bit, byte, and word synchronization.

## **in-band**

Transmission of management protocol over the Fibre Channel.

## **initiator**

A server or workstation on a Fibre Channel network that initiates communications with storage devices. See also [target](#).

## **insistent domain ID mode**

Sets the domain ID of a switch as insistent, so that it remains the same over reboots, power cycles, failovers, and fabric reconfigurations. [interswitch link](#)

See [ISL](#).

## **IOCTL**

I/O control.

## **IP**

Internet Protocol. The addressing part of TCP.

## **ISL**

Interswitch link. A Fibre Channel link from the E\_Port of one switch to the E\_Port of another. See also [E\\_Port](#).

## **isolated E\_Port**

An E\_Port that is online but not operational due to overlapping domain IDs or nonidentical parameters (such as E\_D\_TOVs). See also [E\\_Port](#).

## **IU**

Information unit. A set of information as defined by either an upper-level process protocol definition or upper-level protocol mapping.

## **JBOD**

Just a bunch of disks. A number of disks connected in a single chassis to one or more controllers. See also [RAID](#).

## **K28.5**

A special 10-bit character that indicates the beginning of a transmission word that performs Fibre Channel control and signaling functions. The first seven bits of the character are the comma pattern.

## **key**

A string of data (usually a numeric value) shared between two entities and used to control a cryptographic algorithm. Usually selected from a large pool of possible keys to make unauthorized identification of the key difficult. See also [key pair](#).

## **key pair**

In public key cryptography, a pair of keys consisting of an entity's public and private key. The public key can be publicized, but the private key must be kept secret. See also [public key cryptography](#).

## **L\_Port**

Loop port. A node port (NL\_Port) or fabric port (FL\_Port) that has arbitrated loop capabilities. An L\_Port can be in either fabric mode or loop mode.

## **LAN**

Local area network. A network in which transmissions typically take place over less than 5 kilometers (3.4 miles).

## **latency**

The time required to transmit a frame. Together, latency and bandwidth define the speed and capacity of a link or system.

## **LED**

Light-emitting diode. Indicates the status of elements on a switch.

## **LIFA**

Loop-initialization fabric-assigned frame. Contains a bitmap of all fabric-assigned AL\_PAs and is the first frame transmitted in the loop initialization process after a temporary loop master has been selected.

## **LIHA**

Loop-initialization hard-assigned frame. A hard-assigned AL\_PA that is indicated by a bit set and is the third frame transmitted in the loop initialization process after a temporary loop master has been selected.

## **LILP**

Loop-initialization loop-position frame. The final frame transmitted in a loop initialization process. A returned LIRP contains an accumulation of all of the AL\_PA position maps. This allows loop members to determine their relative loop position. LILP is an optional frame and is not transmitted unless the LIRP is also transmitted.

## **Link Services**

A protocol for link-related actions.

## **LIP**

Loop initialization primitive. The signal that begins initialization in a loop. Indicates either loop failure or node resetting.

## **LIPA**

Loop-initialization previously assigned. The device marks a bit in the bitmap if it did not log in with the fabric in a previous loop initialization.

## **LIRP**

Loop-initialization report position frame. The first frame transmitted in the loop initialization process after all L\_Ports have selected an AL\_PA. The LIRP gets transmitted around the loop so all L\_Ports can report their relative physical position. This is an optional frame.

## **LISA**

Loop-initialization soft-assigned frame. The fourth frame transmitted in the loop initialization process after a temporary loop master has been selected. L\_Ports that have not selected an AL\_PA in a LIFA, LIPA, or LIHA frame select their AL\_PA here.

## **LISM**

Loop-initialization select master frame. The first frame transmitted in the initialization process when L\_Ports select an AL\_PA. LISIM is used to select a temporary loop master or the L\_Port that subsequently starts transmission of the LIFA, LIPA, LIHA, LISA, LIRP, or LILP frames.

## **LM\_TOV**

Loop master time-out value. The minimum time that the loop master waits for a loop initialization sequence to return.

## **loop initialization**

The logical procedure used by an L\_Port to discover its environment. Can be used to assign AL\_PA addresses, detect loop failure, or reset a node.

## **looplet**

A set of devices connected in a loop to a port that is a member of another loop.

## **LR**

Link reset. A primitive sequence used during link initialization between two N\_Ports in point-to-point topology or an N\_Port and an F\_Port in fabric topology. The expected response is an LRR.

## **LRR**

Link reset response. A primitive sequence during link initialization between two N\_Ports in point-to-point topology or an N\_Port and an F\_Port in fabric topology. It is sent in response to an LR and expects a response of Idle.

## **MALLOC**

Memory allocation. Usually applied to buffer credits.

## **MBps**

Megabytes per second.

## **Mbps**

Megabits per second.

## **metric**

A relative value assigned to a route to aid in calculating the shortest path (1000 at 1 Gbps, 500 at 2 Gbps).

## **MIB**

Management Information Base. An SNMP structure to help with device management, providing configuration and device information.

## **MRK**

Mark primitive signal. Used only in arbitrated loop, MRK is transmitted by an L\_Port for synchronization and is vendor specific.

## **MS**

Management Server. Software that allows a storage area network (SAN) management application to retrieve information and administer the fabric and interconnected elements, such as switches, servers, and storage devices. The MS is located at the Fibre Channel well-known address FFFFFAh.

## **multicast**

The transmission of data from a single source to multiple specified N\_Ports (as opposed to all the ports on the network). See also [broadcast](#), [unicast](#).

## **N\_Port**

Node port. A port on a node that can connect to a Fibre Channel port or to another N\_Port in a point-to-point connection. See also [NL\\_Port](#), [Nx\\_Port](#).

## **Name Server**

Simple Name Server (SNS). A switch service that stores names, addresses, and attributes for up to 15 minutes and provides them as required to other devices in the fabric. SNS is defined by Fibre Channel standards and exists at a well-known address. Also called a *directory service*.

## **NL\_Port**

Node loop port. A node port that has arbitrated loop capabilities. Connects an equipment port to the fabric in a loop configuration through an FL\_Port. See also [N\\_Port](#), [Nx\\_Port](#).

## **node**

A Fibre Channel device that contains an N\_Port or NL\_Port. See also [N\\_Port](#), [NL\\_Port](#).

## **node count**

The number of nodes attached to a fabric.

## **node name**

The unique identifier for a node, communicated during login and port discovery.

## **NOS**

Not operational. The NOS primitive sequence is transmitted to indicate that the FC\_Port transmitting the NOS has detected a link failure or is offline and waiting for the offline sequence (OLS) to be received.

## **NS**

Name Server. The service provided by a fabric switch that stores names, addresses, and attributes related to Fibre Channel objects. Can cache information for up to 15 minutes. Also known as the *Simple Name Server* or as a *directory service*. See also [SNS](#).

## **Nx\_Port**

A node port that can operate as either an N\_Port or NL\_Port.

## **OLS**

Primitive sequence offline.

## **ON**

Offline notification. Refers to an ELS field that appears in `portlogdump` command output.

## **OPN**

Open primitive signal. Applies only to arbitrated loops; sent by an L\_Port that has won the arbitration process to open communication with one or more ports on the loop.

## **ordered set**

A transmission word that uses 8B/10B mapping and begins with the K28.5 character. Ordered sets occur outside of frames and include the following items:

- Frame delimiters, which mark frame boundaries and describe frame contents.
- Primitive signals, which indicate events.
- Primitive sequences, which indicate or initiate port states.

Ordered sets differentiate Fibre Channel control information from data frames and manage frame transport.

**originator**

The Nx\_Port that originated an exchange.

**out-of-band**

Transmission of management protocol outside of the Fibre Channel network, usually over Ethernet.

**OX\_ID**

Originator ID. The exchange ID assigned by the originator port.

**parallel**

The simultaneous transmission of data bits over multiple lines.

**path selection**

The selection of a transmission path through the fabric. HP StorageWorks switches use the FSPF protocol. See also [FSPF](#).

**payload**

A Fibre Channel frame has a header and a payload. The payload contains the information being transported by the frame and is determined by the higher-level service or FC\_4 upper-level protocol. There are many different payload formats, based on protocol and size of truck bed. See also [header](#).

**Performance Monitoring**

An HP StorageWorks switch feature that monitors port traffic and includes frame counters, SCSI read monitors, SCSI write monitors, and other types of monitors.

**phantom address**

An AL\_PA value assigned to a device that is not physically in the loop. Also known as *phantom AL\_PA*.

**phantom device**

A device that is not physically in an arbitrated loop but is logically included through the use of a phantom address.

**PID**

Port identifier.

**PKI**

Public key infrastructure. An infrastructure that is based on public key cryptography and certificate authority (CA) and that uses digital certificates. See also [CA](#), [digital certificate](#), [public key cryptography](#).

**PKI certification utility**

Public key infrastructure certification utility. A utility that makes it possible to collect certificate requests from switches and to load certificates to switches. See also [digital certificate](#), [PKI](#).

**PLOGI**

Port login. The port-to-port login process by which initiators establish sessions with targets. See also [FLOGI](#).

**point-to-point**

A Fibre Channel topology that employs direct links between each pair of communicating entities. See also [topology](#).

**port**

In an HP StorageWorks switch environment, an SFP, or GBIC receptacle on a switch to which an optic cable for another device is attached.

**port address**

In Fibre Channel technology, the port address is defined in hexadecimal. In HP Fabric OS, a port address can be defined by a domain and port number combination or by area number. In an ESCON Director, an address used to specify port connectivity parameters and to assign link addresses for attached channels and control units.

**port card**

A hardware component that provides a platform for field-replaceable, hot-swappable ports.

**port log**

A record of all activity on a switch, kept in volatile memory.

**port log dump**

A view of what happens on a switch, from the switch's point of view. The `portlogdump` command is used to read the port log.

**port name**

A user-defined alphanumeric name for a port.

**port swapping**

The ability to redirect a failed port to another port. This feature is available in Fabric OS v4.1.0 and later.

**port\_name**

The unique identifier assigned to a Fibre Channel port. Communicated during login and port discovery.

**POST**

Power-on self-test. A series of tests run by a switch after it is turned on.

**primary FCS switch**

For the HP Secure Fabric OS feature, the primary fabric configuration server switch actively manages security and configurations for all switches in the fabric.

**primitive sequence**

An ordered set that is transmitted repeatedly and continuously. Primitive sequences are transmitted to indicate specific conditions within or conditions encountered by the receiver logic of an FC\_Port. See [OLS](#) and [NOS](#).

**primitive signals**

An ordered set that indicates actions or events and requires just one occurrence to trigger a response. Idle and R\_RDY are used in all three topologies: ARB, OPN, and CLS. MRK is used in arbitrated loop.

**principal switch**

The first switch to boot up in a fabric. Ensures unique domain IDs among roles.

**private key**

The secret half of a key pair. See also [key](#), [key pair](#).

**private loop**

An arbitrated loop that does not include a participating FL\_Port.

**private loop device**

A device that supports a loop and can understand 8-bit addresses, but does not log in to the fabric.

**private NL\_Port**

An NL\_Port that communicates only with other private NL\_Ports in the same loop and does not log in to the fabric.

**protocol**

A defined method and set of standards for communication. Determines the type of error-checking, the data-compression method, how sending devices indicate an end of message, and how receiving devices indicate receipt of a message.

**pstate**

Port State Machine.

**public device**

A device that supports arbitrated loop protocol, can interpret 8-bit addresses, and can log in to the fabric.

**public key**

The public half of a key pair. See also [key](#), [key pair](#).

**public key cryptography**

A type of cryptography that uses a key pair, with the two keys in the pair called at different points in the algorithm. The sender uses the recipient's public key to encrypt the message, and the recipient uses the sender's private key to decrypt it. See also [key pair](#), [PKI](#).

**public loop**

An arbitrated loop that includes a participating FL\_Port and can contain both public and private NL\_Ports.

**public NL\_Port**

An NL\_Port that logs in to the fabric, can function within either a public or a private loop, and can communicate with either private or public NL\_Ports.

**QoS**

Quality of service.

**quad**

A group of four adjacent ports that share a common pool of frame buffers.

**queue**

A mechanism for each AL\_PA address that allows for collecting frames prior to sending them to the loop.

**R\_A\_TOV**

Resource allocation time-out value. The maximum time a frame can be delayed in the fabric and still be delivered. See also [E\\_D\\_TOV](#), [RR\\_TOV](#).

## **R\_CTL**

Route control. The first eight bits of the header, which defines the type of frame and its contents.

## **R\_RDY**

Receiver ready. A primitive signal indicating that the port is ready to receive a frame.

## **R\_T\_TOV**

Receiver transmitter timeout value, used by receiver logic to detect loss of synchronization between transmitters and receivers.

## **RAID**

Redundant array of independent disks. A collection of disk drives that appear as a single volume to the server and are fault tolerant through mirroring or parity checking. See also [JBOD](#).

## **RCS**

Reliable Commit Service.

### **remote switch**

An optional product for long-distance fabrics, requiring a Fibre Channel-to-ATM or SONET gateway.

### **responder**

The N\_Port with which an exchange originator wishes to communicate.

## **RLS**

Read Link Status.

### **route**

As applied to a fabric, the communication path between two switches. May also apply to the specific path taken by an individual frame, from source to destination. See also [FSPF](#).

### **routing**

The assignment of frames to specific switch ports, according to frame destination.

## **RR\_TOV**

Resource recovery timeout value. The minimum time a target device in a loop waits after a LIP before logging out an SCSI initiator. See also [E\\_D\\_TOV](#), [R\\_A\\_TOV](#).

## **RSCN**

Registered state change notification. A switch function that allows notification of fabric changes to be sent from the switch to specified nodes. The fabric controller issues RSCN requests to N\_Ports and NL\_Ports, but only if they have registered to be notified of state changes in other N\_Ports and NL\_Ports. This registration is performed via the State Change Registration (SCR) Extended Link Service. An N\_Port or NL\_Port can issue an RSCN to the fabric controller without having completed SCR with the fabric controller.

## **RTWR**

Reliable transport with response. May appear as a task in `portlogdump` command output.

### **running disparity**

A binary parameter that indicates the cumulative disparity (positive or negative) of all previously issued transmission characters.

## **RW**

Read/write. Refers to access rights.

## **RX**

Receiving frames.

## **RX\_ID**

Responder exchange identifier. A 2-byte field in the frame header that can be used by the responder of the exchange to identify frames as being part of a particular exchange.

## **S\_ID**

Source ID. Refers to the native port address (24 bit address).

## **SAN**

Storage area network. A network of systems and storage devices that communicate using Fibre Channel protocols. See also [fabric](#).

## **SAN architecture**

The overall design of a storage network solution, which includes one or more related fabrics, each of which has a topology.

## **SAN port count**

The number of ports available for connection by nodes in the entire SAN.

## **SCN**

State change notification. Used for internal state change notifications, not external changes. This is the switch logging that the port is online or is an Fx\_port, not what is sent from the switch to the Nx\_ports.

## **SCR**

State change registration. Extended Link Service (ELS) requests the fabric controller to add the N\_Port or NL\_Port to the list of N\_Ports and NL\_Ports registered to receive the Registered State Change Notification (RSCN) Extended Link Service.

## **SCSI**

Small Computer Systems Interface. A parallel bus architecture and a protocol for transmitting large data blocks to a distance of 15 to 25 meters.

## **secelnet**

A protocol similar to telnet but with encrypted passwords for increased security.

## **Secure Fabric OS**

A separately sold HP feature that provides advanced, centralized security for a fabric.

## **security policy**

Rules that determine how security is implemented in a fabric. Security policies can be customized through HP Secure Fabric OS or HP Fabric Manager.

## **SEQ\_ID**

Sequence identifier. A one-byte field in the frame header change to identify the frames as being part of a particular exchange sequence between a pair of ports.

## **sequence**

A group of related frames transmitted in the same direction between two N\_Ports.

## **sequence initiator**

The N\_Port that begins a new sequence and transmits frames to another N\_Port.

## **sequence recipient**

Serializing and deserializing circuitry. A circuit that converts a serial bit stream into parallel characters, and vice-versa.

## **SFP**

Small-form-factor pluggable. A transceiver used on 2 Gbps or 4 Gbps switches that replaces the GBIC.

## **Simple Name Server (SNS)**

A switch service that stores names, addresses, and attributes for up to 15 minutes and provides them as required to other devices in the fabric. SNS is defined by Fibre Channel standards and exists at a well-known address. Also called a *directory service* or *Name Server*.

## **Single CP Mode**

The -s option of the `firmwaredownload` command. Using `firmwaredownload -s` enables Single CP Mode. In the Core Switch 2/64, Single CP Mode enables a user to upgrade a single CP card and to select full-install, auto-reboot, and auto-commit.

## **SNMP**

Simple Network Management Protocol. An Internet management protocol that uses either IP for network-level functions and UDP for transport-level functions, or TCP/IP for both. Can be made available over other protocols, such as UDP/IP, because it does not rely on the underlying communication protocols.

## **SNS**

See [Simple Name Server \(SNS\)](#).

## **SOF**

Start of frame. A group of ordered sets that marks the beginning of a frame and indicates the class of service the frame will use.

## **SONET**

Synchronous optical network. A standard for optical networks that provides building blocks and flexible payload mappings.

## **special character**

A 10-bit character that does not have a corresponding 8-bit value but is still considered valid. The special character indicates that a specific transmission word is an ordered set. This is the only type of character to have five 1s or 0s in a row.

## **SSH**

Secure shell. Used starting in HP Fabric OS v4.1.0 to support encrypted telnet sessions to the switch. SSH encrypts all messages, including the client sending the password at login.

## **switch**

A fabric device providing bandwidth and high-speed routing of data through link-level addressing.

### **switch name**

The arbitrary name assigned to a switch.

### **switch port**

A port on a switch. Switch ports can be E\_Ports, F\_Ports, or FL\_Ports. See also [E\\_Port](#), [F\\_Port](#), [FL\\_Port](#).

### **switch-to-switch authentication**

The process of authenticating both switches in a switch-to-switch connection using digital certificates. See also [authentication](#), [digital certificate](#).

## **syslog**

Syslog daemon. Used to forward error messages.

## **T11**

A standards committee chartered with creating standards for Fibre Channel.

### **target**

A storage device on a Fibre Channel network. See also [initiator](#).

## **TC**

Track Changes.

### **telnet**

A virtual terminal emulation used with TCP/IP. Telnet is sometimes used as a synonym for the HP Fabric OS CLI.

### **tenancy**

The time beginning when a port wins arbitration in a loop until the same port returns to the monitoring state. Also called *loop tenancy*.

### **Time Server**

A Fibre Channel service that allows for the management of all timers.

### **topology**

As applied to Fibre Channel technology, the configuration of the Fibre Channel network and the resulting communication paths allowed. There are three possible topologies:

- Point to point, which is a direct link between two communication ports.
- Switched fabric: multiple N\_Ports linked to a switch by F\_Ports.
- Arbitrated loop: multiple NL\_Ports connected in a loop.

## **Track Changes**

An HP Fabric OS feature that can be enabled to report specific activities (for example, logins, logouts, and configuration task changes). The output from the track-changes feature is dumped to the system message log for the switch.

## **transceiver**

A device that converts one form of signaling to another for transmission and reception; in fiber optics, optical to electrical.

## **translative mode**

A mode in which private devices can communicate with public devices across the fabric.

## **transmission character**

A 10-bit character encoded according to the rules of the 8B/10B algorithm.

## **transmission word**

A group of four transmission characters.

## **trap (SNMP)**

The message sent by an SNMP agent to inform the SNMP management station of a critical error. See also [SNMP](#).

## **trunking**

In Fibre Channel technology, a feature that enables distribution of traffic over the combined bandwidth of up to four ISLs between adjacent switches, while preserving in-order delivery.

## **trunking group**

A set of up to four trunked ISLs.

## **trunking ports**

The ports in a set of trunked ISLs.

## **TS**

Time Server.

## **TTL**

Time-to-live. The number of seconds an entry exists in cache before it expires.

## **tunneling**

A technique for enabling two networks to communicate when the source and destination hosts are both on the same type of network but are connected by a different type of network.

## **TX**

Transmit.

## **U\_Port**

Universal port. A switch port that can operate as a G\_Port, E\_Port, F\_Port, or FL\_Port. A port is defined as a U\_Port when it is not connected or has not yet assumed a specific function in the fabric.

## **unicast**

The transmission of data from a single source to a single destination. See also [broadcast](#), [multicast](#).

## **UTC**

Universal Time Conversion. Also known as *Coordinated Universal Time*, which is an international standard. UTC is 8 hours later than Pacific Standard Time and 5 hours later than Eastern Standard Time. See also [GMT](#).

## **WAN**

Wide area network.

## **watchdog**

A software daemon that monitors Fabric OS modules on the kernel.

## **well-known address**

For Fibre Channel technology, a logical address defined by Fibre Channel standards as assigned to a specific function and stored on the switch.

## **WWN**

World Wide Name. An identifier that is unique worldwide. Each entity in a fabric has a separate WWN.

## **zone**

A set of devices and hosts attached to the same fabric and configured as being in the same zone. Devices and hosts within the same zone have access to others in the zone, but are not visible to any outside the zone.

## **zone configuration**

A specified set of zones. Enabling a configuration enables all zones in that configuration. See also [defined zone configuration](#).

## **zoning**

A feature in fabric switches or hubs that allows segmentation of a node by physical port, name, or address.



# Index

## A

account ID 19  
account privilege levels 21  
activating  
  a switch certificate 58  
adding  
  end-to-end monitors 146  
  filter-based monitors 150  
adding members  
  zone 133  
adding members to a zone configuration 135  
adding members, alias 132  
adding switches to a zone 137  
Advanced Performance Monitoring commands 143  
alias  
  adding members 132  
  creating 131  
  deleting 132  
  removing members 132  
ATM gateway 225  
authorized reseller, HP 11

## B

beaconing mode 98  
boot PROM password 70, 72  
browser and Java support 54  
buffer-limited port 107

## C

chassisshow command 34  
choosing  
  a CA 55  
clearing  
  monitor counters 155  
command  
  chassisshow 34  
  configupload 77  
  fabricshow 34  
  hashow 34  
  licenseadd 26  
  licenseremove 26  
  licenseshow 26  
  nsallshow 34  
  nsshow 34  
  slotshow 34  
  switchshow 33  
commands  
  Advanced Performance Monitoring 143  
configupload command 77  
configuration  
  save to a host 75

configuring for the SSL protocol 54  
configuring the browser 58  
configuring the radius server 48  
configuring the switch 51  
connection  
  serial 19, 20  
  telnet 19  
conventions  
  document 10  
  text symbols 10  
core/edge topology and ISL trunking 114  
CRC errors, displaying 145  
creating a zone 133  
creating a zone configuration 134  
creating an alias 131

## D

date and time 22  
default names 27  
defined zone configuration 126  
deleting  
  end-to-end monitors 149  
  filter-based monitors 152  
deleting a zone 134  
deleting a zone configuration 135  
deleting an alias 132  
designing fabric for trunking 114  
deskew values  
  displaying 119  
device  
  connecting 31  
device-based routing 99, 100  
disabled zone configuration 127  
disabling interoperability mode 223  
displaying  
  CRC error count 145  
  end-to-end mask 148  
displaying and deleting certificates 59  
document  
  conventions 10  
  prerequisites 9  
  related documentation 9  
domain ID  
  FICON environment 31

## E

enabling a card 91  
enabling and disabling local authentication 53  
enabling interoperability mode 223  
encryption 54  
end-to-end monitoring 145

end-to-end monitors  
  adding 146  
  deleting 149  
  displaying the mask 148  
  restoring configuration 156  
  saving configuration 156  
  setting a mask 148  
example  
  chassisshow 34  
  fabricshow 34  
  nsallshow 34  
  slotshow 34  
exchange-based routing 99, 100

**F**  
fabric connectivity 34  
fabric, designing for trunking 114  
fabricshow command 34  
fans, status of 171  
feature licenses 25  
FICON environment  
  changing domain ID 31  
filter-based monitoring 149  
filter-based monitors  
  adding 150  
  deleting 152  
  restoring configuration 156  
  saving configuration 156  
frame transfer with Brocade Remote Switch 225

**G**  
gateway, Remote Switch 225  
generating  
  batch of licenses 25  
generating a public/private key 55  
generating and storing a csr 56  
getting help 11

**H**  
hard zoning 127  
hardware-enforced zoning 127  
hashow command 34  
high availability (HA) 34  
host-based zoning 124  
HP  
  authorized reseller 11  
  storage web site 11  
  technical support 11

**I**  
installing a root certificate to the Java plug-in 59  
installing a switch certificate 57  
interswitch link 31  
ISL 31

**J**  
Java support, SSL 54

**K**  
key  
  license, generating on the web 25  
  transaction, for licensed features 25

**L**  
license key  
  activating 26  
license keys  
  generating 25  
licenseadd command 26  
licensed features 25  
licenseremove command 26  
licenseshow command 26  
login  
  switch 19  
long distance ISLs 108  
LUN masking 124

**M**  
mask for end-to-end monitors  
  displaying 148  
  setting 148  
monitors  
  clearing counters 155

**N**  
name server zoning 124  
network security 39  
nsallshow command 34  
nsshow command 34

**O**  
obtaining certificates 57

**P**  
password 19  
  boot PROM 70  
passwords  
  recovering forgotten passwords 73  
perfAddEEMonitor command 146  
perfAddIPMonitor command 150  
perfAddUserMonitor command 150  
perfCfgRestore command 156  
perfCfgSave command 156  
perfDelEEMonitor command 149  
perfDelFilterMonitor command 152  
performance monitoring commands 143  
perfSetPortEEMask command 148  
perfShowAlpaCRC command 145  
perfShowPortEEMask command 148  
PKI 54  
port  
  buffer-limited 107  
port-based routing 99  
ports  
  status of 168  
powering off a card 90

prerequisites 9  
privileges in accounts 21  
public key infrastructure encryption 54

## R

recovering forgotten passwords 73  
recovery password 72  
recovery string, boot PROM password 70  
related documentation 9  
Remote Switch 225  
removing  
  end-to-end monitors 149  
  filter-based monitors 152  
removing members  
  zone 133  
removing members from a zone configuration 135  
removing members, alias 132  
restoring monitor configuration 156  
restoring the system configuration settings 76  
routing  
  assigning a static route 100  
routing policies 99, 100

## S

saved zone configuration 127  
saving monitor configuration 156  
secure shell (SSH) 39  
secure socket layer protocol 54  
secure sockets layer 54  
security 39  
  security and zoning 139  
serial connection 19, 20  
setting mask for end-to-end monitors 148  
setting the switch date and time 22  
slotshow command 34  
SSL 54  
standard filter-based monitors 150  
static route 100  
storage-based zoning 124  
summary of SSL procedures 54  
supportsave command 177  
switch  
  system status 166  
switch names 27  
switchshow command 33  
symbols  
  in text 10  
symbols in text 10

## T

technical support, HP 11  
telnet connection 19  
temperature, status of 172  
text symbols 10  
time and date 22  
traffic patterns  
  planning for 114  
transaction key 25

troubleshooting certificates 60  
trunking  
  displaying information 119

## U

user-defined filter-based monitors 150

## V

verify  
  device connectivity 31  
  fabric connectivity 34  
  high availability (HA) 34  
viewing  
  fan status 171  
  port status 168  
  power supply status 172  
  temperature status 172  
viewing an alias  
  alias  
    viewing 132  
viewing zone database configurations 136  
viewing zones 134

## W

websites, HP storage 81, 84

## Z

zone  
  adding members 133  
  adding switches 137  
  creating 133  
  creating a configuration 134  
  deleting 134  
  deleting a configuration 135  
  removing members 133  
  viewing 134  
  viewing configurations 136  
zone aliases 126  
zone configuration  
  adding members 135  
  removing members 135  
zone configurations 126  
zone objects 126  
zoning  
  administering security 139  
zoning schemes 127

